Pricing and Online Platform Design

by

Titing Cui

B.S., Zhejiang University, 2017

M.S., KTH Royal Institute of Technology, 2019

Submitted to the Graduate Faculty of

the Joseph M. Katz Graduate School of Business in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2024

UNIVERSITY OF PITTSBURGH

JOSEPH M. KATZ GRADUATE SCHOOL OF BUSINESS

This dissertation was presented

by

Titing Cui

It was defended on

June 17th 2017

and approved by

Michael L. Hamilton (Chair), PhD, Assistant Professor

Joseph M. Katz Graduate School of Business, University of Pittsburgh

Prakash Mirchandani (Co-Chair), PhD, Professor

Joseph M. Katz Graduate School of Business, University of Pittsburgh

Jennifer Shang, PhD, Professor

Joseph M. Katz Graduate School of Business, University of Pittsburgh

Esther Gal-Or, PhD, Professor

Joseph M. Katz Graduate School of Business, University of Pittsburgh

Sera Linardi, PhD, Associate Professor

Graduate School of Public and International Affairs, University of Pittsburgh

Pricing and Online Platform Design

Titing Cui, PhD

University of Pittsburgh, 2024

Revenue management and pricing is an evolving discipline that is rapidly expanding with the rise of online marketplaces. For many online marketplaces, pricing strategy is a central part of the platform's design. In this thesis we employ methods from operations research, stochastic optimization, and economics to model online platforms and optimize their operations. Throughout we highlight the unique operational characteristics of these online platforms, and how they require new and sophisticated models to study them. Specifically, this dissertation focuses on the study of emerging pricing and online platform design problems in three contexts. First, we consider the problem of market segmentation and pricing under the assumption that the seller has trained a regression model that maps customer features to valuations. Second, we study the pricing strategy for online dating platforms, where we analyze the profit and welfare trade-offs associated with different length subscriptions. Third and finally, we consider the design and optimization of rating system for platforms, focusing on how to extend rating systems to mitigate issues resulting from stale reviews.

Table of Contents

List of Tables

List of Figures

Preface

First and foremost, I would like to thank my advisors, Professor Michael L Hamilton, and Professor Prakash Mirchandani. Prakash guided me through the first two years of my PhD life, especially during COVID-19. The courses we picked together prepared me to be capable of my following research. As my thesis advisor, Mike introduced me to the field of revenue management and pricing. Mike taught me how to write clean proofs, how to make insightful slides, and how to prepare for my job market presentation. I enjoyed the time we spent in your office, talking about research ideas, proof details, and funny jokes. This dissertation would not have been possible without your continuous guidance, support, and encouragement.

I would like to thank my all coauthors, especially Su, Thomas, and Mike. Thank you for offering the opportunity to work with you, and sharing the job market experience with me. Thank you for correcting my bad behavior in formal mathematical writing and coding.

I would like to thank all the other members of my committee, Professor Jennifer Shang, Professor Esther Gal-or, and Professor Sera Linardi. Thank you all for serving on my dissertation committee, and for your invaluable comments during my proposal. I would also like to thank other professors at Katz, especially Professor Arian Aflaki, Professor Leon Valdes, and Professor Pandu Tadikamalla. Thank you for your help during my PhD study and your advice when I was on the job market. I have learned a lot from Leon's lecture notes for my teaching.

I would also like to thank all my friends that I made and all the people that I met throughout my PhD study. Especially my office mates, we spent so much time together, and also the PhD students at Tepper OM, thank you for inviting me to your OM nights, YinzOR conference.

My thanks also go to staff members at the PhD office, Carrie, Melissa, and Rachael. Thank you for reminding the timeline of my PhD study, the reimbursement for my academic travels, and always welcoming me with big smiles.

I would like to thank all my families, for always being supportive during my journey. I

am immensely grateful for the love and support I have received from them. Lastly, I would like to thank my wonderful partner Huanhuan, who has always believed in me.

To my parents, for everything

1.0 Introduction

The emergence of online platforms has redefined how markets function, created new economic opportunities, and driven significant changes in consumer behavior and business operations. One of the key challenges for online platforms is achieving a balance between attracting a sufficient user base and monetizing services effectively. Many online marketplaces must carefully design their pricing structures to ensure accessibility while also incentivizing premium features or subscriptions. In this thesis, we study three emerging topics, inspired by modern online marketplaces at the intersection of pricing and platform design.

1.0.1 Summaries and Contributions

In this section, we summarize the main contributions of Chapters 2, 3, and 4. Each of the following Chapters will be a self-contained exposition on a single emerging topic in online platform design and pricing.

In Chapter 1, we explore semi-personalized pricing strategies where a seller uses customer features to segment the market and offer segment-specific prices. Typically, finding optimal market segmentation and pricing policies is computationally challenging, leading practitioners to rely on heuristic methods. We address this issue by studying the optimization and analysis of feature-based market segmentation and pricing strategies, under the assumption that the seller has trained a regression model that maps customer features to valuations. In this framework, we present novel results on the hardness and approximation of these strategies. When model noise is independent and log-concave, we demonstrate that the joint segmentation and pricing problem can be solved efficiently. We also analyze structural properties of the optimal feature-based market segmentation and pricing. Finally, we conduct a case study using home mortgage data and show that our approach can achieve nearly all of the available revenue with only a few segments, significantly outperforming heuristic methods. Additionally, we provide insights into the structural properties of pricing from regression models that may be of independent interest. This chapter is based on joint work with Professor Michael Hamilton.

In Chapter 2, we develop a novel model for the operation of an online dating platform. Online dating is the most common way for new US couples to meet, with more than three-in-ten Americans having used dating apps, and with revenues from dating apps surpassing five billion dollars annually. Most of these dating apps generate revenue through subscription-based pricing, where subscriptions for a fixed period of app access are sold at recurring prices. Subscription-based pricing is a common monetization strategy for mobile apps, but in the context of online dating, different subscription period lengths can lead to very different outcomes for the platform and its users. In this study we analyze the profit and welfare trade-offs associated with different length subscription pricing strategies for online dating platforms. By analyzing our model we find that short period pricing always achieves at least $\frac{1}{1+e} \approx 26.9\%$ of the profit of the optimal policy. However when the marginal cost is low, longer period pricing yields optimal revenue and higher welfare. Longer period pricing also provably coordinates incentives between users and platforms, allowing the platform to incorporate user preference information in a way that leads to lower prices. This chapter is based on joint work with Professor Michael Hamilton.

In Chapter 3, we model a two-sided ecosystem where an online platform facilitates interactions between workers and customers. Specifically, we consider a worker who, when hired, performs a service on a platform for which they receive a rating. We focus on how the rating system should process the sequence of worker outcomes, including the null outcome where the worker does not transact on the platform. To mitigate issues due to sequential staling, we introduce a new class of α-moving average rating systems which systematically down-weight older reviews. We show that a moving average approach generates ratings in a way which optimally protects customers from changes in the underlying quality of service that may otherwise go undetected due to staling. To address issues due to temporal staling, we extend our rating system to include penalty reviews that are incurred by service providers after long periods of inactivity. In a natural model of online platforms that facilitate interactions between service providers and customers, we analyze how to set the parameters and the penalty term for a rating system so as to maximize platform revenue. We give a prior-free way of choosing these parameters which approximately maximizes revenue and also mitigates temporal staling by encouraging service providers to continuously solicit new reviews. This chapter is based on joint work with Professor Michael Hamilton.

2.0 Optimal Feature-Based Market Segmentation and Pricing

2.1 Introduction

Third-degree price discrimination concerns the now ubiquitous practice of selling a good at different prices for different consumers [131]. For example, in the sale of proprietary software licenses, prices may differ based on whether the customer is a student, professional, or corporate user (see Section 2.1 for an example and [86] for an extended discussion). In insurance markets, firms gather extremely rich and nuanced feature information about their customers, ostensibly to estimate risk, but which is also leveraged to price discrimination on the basis of demographic and/or geographic information [28, 130]. In commercial markets facing walk-in customers, firms are comparatively limited in their information about their customers but can still profitably leverage feature information to price discriminate, for instance in movie theater ticket sales the customer's age (child, adult, senior) and the time of screening (weekend, matinee) can be used to issue semi-personalized prices via discounts [45].

Each of these markets vary in both the quality and descriptive power of the information they have about their customers, as well as in the operational difficulty of setting and changing their prices, and thus implementing price discrimination. When the information about customers in the market is of low quality or is largely censored, fully personalized price discrimination where each customer is charged a personalized price may be futile, however as mentioned above, that does not preclude the use of some modest price discrimination via market segmentation. In fact, even when information is richly textured and pricing is largely unconstrained by legal and/or operational considerations, a small static set of prices based on customer features is still often preferred to fully personalized pricing. A small set of market segments and prices is conceptually and operationally simple to implement, and a surprisingly small number of options is often sufficient to achieve strong revenue [41]. In this chapter, we propose a general framework for studying semi-personalized pricing strategies which can capture these variations in predictive power and operational flexibility, which we term *feature-based market segmentation and pricing* (FBMSP).

Figure 1: Example of feature-based market segmentation and pricing.

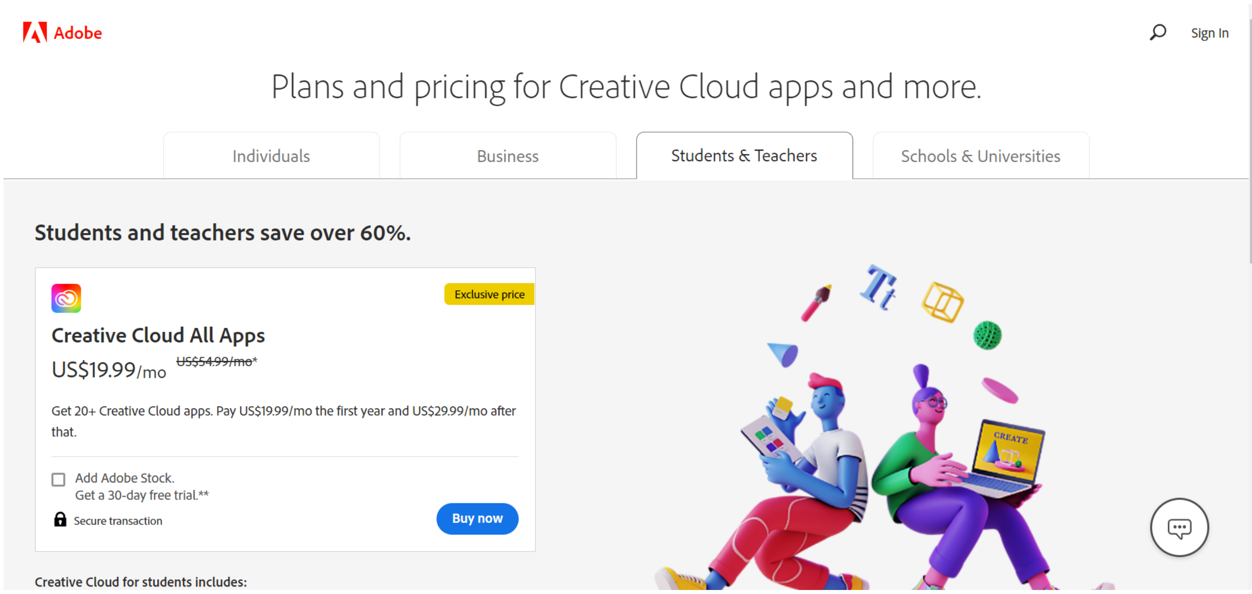

Note. An example of feature-based market segmentation and pricing for Adobe Creative Cloud products (see `https://www.adobe.com/creativecloud/plans.html`). Here customers are segmented based on their attributes (i.e. student versus professional) and prices vary based on the segment customers are in, with the student rate being 60% of the professional rate (note discounts are enforced by requiring a valid .edu email). Further, note the number of segments, k, is only four. It will be informative to think of k as ≈ 4 in this work.

Finding and optimizing generic market segmentation and pricing policies is a well-studied problem in industry with academic roots in operations research/management, marketing, economics, and computer science. However, archetypal formulations of the segmentation and pricing problem are well known to be intractably hard [81, 82]. To deal with this hardness, much of the literature has taken a heuristic approach to the problem [38, 12, 33, 89, 87], separating the segmentation and pricing components. Segmentation-then-price procedures use tools from unsupervised learning to first identify consumer segments/clusters with similar features, and then identify revenue optimal prices for those chosen market segments. Using our framework, we advance the study of segmentation and pricing by finding *jointly optimal* segmentation and pricing under some realistic assumptions about how firms leverage feature information to predict customers' valuations. Specifically, in practice, a firm's valuation model i.e. the model that maps features to valuation or a proxy for willingness-to-pay, is built using regression. Regression models come with their own theory and standard set of assumptions that we profitably utilize to study market segmentation and pricing as well. We show that by leveraging the assumptions of independence and log-concavity of residuals in the regression model, the resulting revenue-maximizing feature-based market segmentation and pricing (FBMSP) enjoys a simple, intuitive structure, and can be computed efficiently. Further, our structural results allow us to analyze optimal FBMSP and derive new managerial insights about such policies, including guidance for choosing the number of segments, and conditions for when segmentation and pricing are near-optimal.

2.1.1 Our Contributions

To summarize our contributions:

1. We first study the algorithmic problem of finding the optimal FBMSP. In general, the problem is intractably hard, so we focus our attention on the case when valuations are predicted according to a regression model with independent residuals. We show that with no additional assumptions, while we can prove some promising structural properties (c.f. Lemma 1) and provide a $(1-1/e)$ approximation algorithm for the optimal segmentation and pricing (c.f. Remark 1), unfortunately finding the optimal FBMSP is still NP-Hard to compute (c.f. Theorem 1). However, when we further assume the residuals are log-concave, as is often the case, we are able to evade our hardness result. Specifically, when residuals are independent and log-concave, we prove the optimal policy has a simple *interval* structure which allows us to compute in it quadratic time via dynamic programming (c.f. Theorem 2).
2. We next turn our attention to analyzing the performance of optimal feature-based market segmentation and pricing. Specifically, we consider the practical operational question of how to choose the number of segments k so as to guarantee minimal loss against a fully personalized pricing benchmark. We show three results that can help guide practitioners in choosing k. First, we show that an upper bound on the loss against personalized pricing

can be achieved by simply examining the loss in the model, ignoring the noise term (c.f. Theorem 3). Second, we tightly upper bound the optimal rate at which FBMSP tends to personalized pricing as a function of the number of segments k and some valuation parameters (c.f. Theorem 4). Finally, we show that the revenue of FBMSP is concave in k (c.f. Theorem 5). Taken together, these three results allow a practitioner to use their regression model (without reference to the complicating error!) to find k via a simple elbow method, and feel confident that the results of such a heuristic are provably close to optimal.

3. Finally in Section 2.5, we demonstrate our method on real housing loan data collected in Pennsylvania in 2020, and compare its performance against standard segment-then-price methodologies. We find our approach significantly outperforms heuristic methods, especially when the number of segments is small and the variation in the valuations comes primarily from variation in the regression model $\mu(\cdot)$, as opposed to variation from the prediction error ϵ. We also note that the segmentations found by our approach are qualitatively different than those in segment-then-price, with our approach quickly isolating key differences between groups, whereas heuristic approaches can get bogged down in pointless price discrimination between groups until it discovers the important differences for the revenue.

2.1.2 Literature Review

Our work is influenced by, and contributes to, several streams of literature across operations management, marketing, and computer science. We now overview some of these streams and connect them to our work.

2.1.2.1 Theory of Price Discrimination

There is extensive literature on the theory of pricing discrimination beginning in economics, and spanning operations management, marketing, and computer science. Much of the classic literature in this area ([113, 99, 80, 132, 119, 20, 42, 137]) focuses on the impact of price discrimination on social welfare, or the effects of price discrimination on the resultant equilibrium prices. In this chapter, we investigate market segmentation and pricing from the perspective of a revenue-maximizing monopolist, focusing on computational/practical implementations of such policies.

Specifically, in the language of [132] we study third-degree price discrimination which concerns when a company charges a different price to different consumer groups. In practice, third-degree price discrimination is the most common form of price discrimination, with companies leveraging additional information about consumer features to offer different prices to different implicit/explicit segments in a variety of ways ([123, 75, 22, 36, 39, 53]). Several papers have analyzed the value of such price discrimination tactics compared to uniform pricing ([70, 52]). In contrast, we investigate the value of the optimal feature-based market segmentation and pricing in this chapter, and compare this type of semi-personalized pricing against a fully personalized benchmark.

2.1.2.2 Regression Based Price Discrimination

In recent years, data-driven pricing strategies have become increasingly common ([35, 56, 120, 8, 50, 102, 24, 51]). In these works, customers are offered a personalized price based on features that are predictive of their valuation of the product, especially by tree-based prescriptive approaches ([13, 77, 21, 24]). Unlike most data-driven pricing literature, in our work, we ignore how the regression model is found and instead take the prediction of customer's valuation as input, and analyze how it may be profitable leveraged to compute and analyze optimal FBMSP.

2.1.2.3 Algorithms for Market Segmentation and Pricing

This chapter contributes to a line of literature studying market segmentation and pricing from an algorithmic/computational complexity perspective. Indeed many models of joint market segmentation and pricing are known to be intractably hard to compute going back at least to the pioneering work of [81, 82], restricting their applicability in practice. Often in marketing, to evade these hardness results the segmentation and pricing decisions are made sequentially instead of being evaluated together ([44]), and at first blush it seems that Theorem 1 implies our model, for all the structure gained through independence, is ultimately no better. Fortunately, we will see for almost all regression models in practice our model makes jointly optimal segmentation and pricing tractable and well structured.

If the regression error is log-concave, as we assume in Section 2.3.2, computing the optimal feature-based segmentation is structurally similar to the *1D Clustering* problem for which dynamic programming approaches have been employed (see [62] for a modern overview), and can be solved in polynomial time. Other algorithmic approaches for feature-based pricing can be seen in [40, 105, 74], albeit in different models.

2.1.3 Chapter Outline

The remainder of this chapter is organized as follows. In Section 2.2 we introduce our model for FBMSP and provide some preliminary structural results. In Section 2.3 we study the problem of computing the revenue-optimal FMBSP. In Section 2.4 we analyze the structure of revenue-optimal FBMSP and provide some theory to guide practitioners in choosing the number of segments, k. In Section 2.5 we demonstrate our approach on a well known Home Mortgage Disclosure Act dataset. Finally, in Section 2.6 we provide concluding remarks and highlight future directions. All examples and proofs referenced in the main body can be found in Appendices A.1 and A.2 in the Appendix.

2.2 Model and Preliminaries

We consider a revenue-maximizing seller offering a good in unlimited supply. For simplicity of presentation, we will assume the good is produced costlessly and so revenue and profit are equivalent (we note the model presented in this chapter and all results easily extend to the case when each good has a per unit cost c). We further assume each customer in the market is described by some feature vector $\mathbf{x}$ of their observable characteristics, and has some valuation for the good which depends on their feature vector, $V|\mathbf{x}$. The market characteristics as a whole can be described as a distribution over the feature vectors $\mathbb{X} \sim F_{\mathbb{X}}$, which is supported on some feature space $\mathcal{X} := \text{supp}(\mathbb{X})$. These features vectors can consist of any information about the customers, including demographic information like gender, household status, income etc.

In line with modern practice, we model the seller as having trained some regression model $\mu : \mathcal{X} \to \mathbb{R}^+$ to predict a customer's valuation for a good from their feature vector. We assume the regression model has residual error ϵ but is correct in expectation, so that the predicted valuation for a customer with features $\mathbf{x}$ is $\mu(\mathbf{x}) := \mathbb{E}[V|\mathbf{x}]$, and the valuation model is $V = \mu(\mathbb{X}) + \epsilon$. We will use F to be the distribution of the valuations V, $F_{\mathbb{X}}$ to be the distribution of the feature vectors, F_ϵ to be the distribution of the error term ϵ, and $f_{\mathbb{X}}$, f_ϵ, and f to be the densities, respectively. We will use $\overline{\mathrm{F}}$ to denote the survival function, i.e., $\overline{\mathrm{F}}(x) := 1 - F(x)$.

Figure 2: An example of FBMSP.

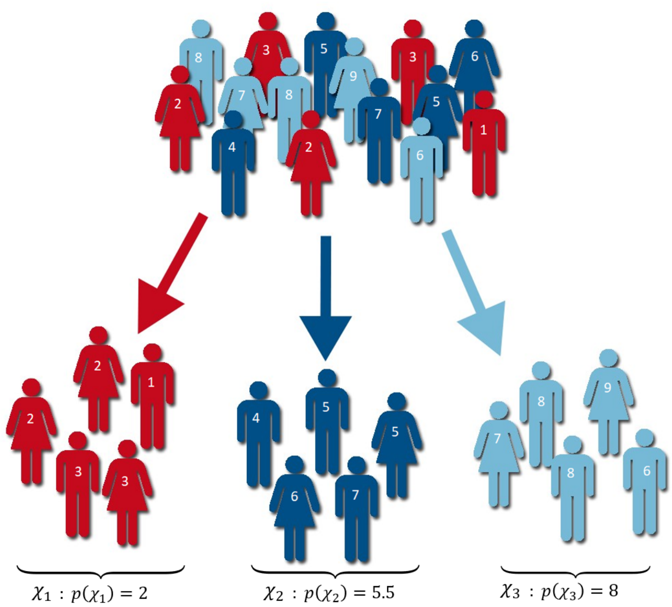

Note. Depicted is an example of feature-based market segmentation. For each customer some numeric prediction of their valuation is given. The feature space $\mathcal{X}$ consists of all combinations of the color and gender for the customer, and the depicted feature-based market segmentation leverages color (not necessarily optimally) to sort them into three segments $\mathcal{X}_i$, $i \in [3]$, each with a distinct segment level price, $p(\mathcal{X}_i)$.

For a seller with a valuation model $\mu(\cdot)$, we will study the revenue achievable by selling strategies where the feature space of the market, $\mathcal{X}$, is partitioned into k segments $\{\mathcal{X}_i\}_{i=1}^k$, $\cap\mathcal{X}_i = \emptyset$, $\cup\mathcal{X}_i = \mathcal{X}$, such that on each segment the seller offers a distinct price $p(\mathcal{X}_i)$. Now we are ready to define feature-based market segmentation and pricing strategies, which is the main object of this study.

Feature-Based Market Segmentation and Pricing: In feature-based market segmentation and pricing the seller partitions the feature space into k segments $\{\mathcal{X}_i\}_{i=1}^k$, and on each segment offers a single price $p(\mathcal{X}_i)$ (see Section 2.2 for example). The expected profit of such a segmentation is,

$$\mathcal{R}_{kXP}\left(\{\mathcal{X}_i\}_{i=1}^k, \{p(\mathcal{X}_i)\}_{i=1}^k\right) := \sum_{i=1}^k p(\mathcal{X}_i) \int_{\mathbf{x}\in\mathcal{X}_i} \Pr\left(\mu\left(\mathbf{x}\right) + \epsilon \geq p(\mathcal{X}_i)\right) f_{\mathbb{X}}(\mathbf{x}) d\mathbf{x}, \tag{1}$$

where the sum is over the k market segments, and the revenue of each segment is the segment price $p(\mathcal{X}_i)$ times the probability of sale at the price, integrated over the feature vectors in the segment. Given a segmentation it will often be convenient to think of the prices as the revenue optimal ones for that segment. To that end, we denote the optimal price on segment $\mathcal{X}_i$ by $p_\epsilon(\mathcal{X}_i)$ i.e.,

$$p_\epsilon(\mathcal{X}_i) := \arg\max_p p \int_{\mathbf{x}\in\mathcal{X}_i} \Pr\left(\mu\left(\mathbf{x}\right) + \epsilon \geq p\right) f_{\mathbb{X}}(\mathbf{x}) d\mathbf{x}.$$

We will use $\mathcal{R}_{kXP} := \max_{\mathcal{X}_1,\ldots,\mathcal{X}_k} \sum_{i=1}^k \mathcal{R}_{kXP}(\{\mathcal{X}_i\}_{i=1}^k, \{p_\epsilon(\mathcal{X}_i)\}_{i=1}^k)$ denote the optimal profit for a feature-based (k) market segmentation and pricing strategy.

Note that our framework for feature-based market segmentation and pricing is very flexible, and captures many well studied models as special cases. For instance when $k = 1$, FBMSP is just the revenue generated by a single static price for the good (sometimes referred to as the revenue of the monopoly price, or posted price, or single price). Similarly when the number of segments is very large, i.e. $k \to \infty$, FBMSP becomes the revenue of feature-based personalized pricing, where each customer receives an appropriately personalized price. The revenue of feature-based personalized pricing is a useful upper bound to compare with the revenue of optimal FBMSP for some fixed number of segments k. In Section 2.4 we consider the question of how large does k need to be, in general, to approach the revenue of personalized pricing, with the hope that a reasonably small k should suffice (see [52] for a detailed discussion of when feature-based personalized pricing is provably close or far from the revenue of a single price).

Moreover, when the error distribution ϵ is 0 almost surely (a.s.), our model represents the achievable revenue in the world of the prediction model μ without regard to the models potential error. We term this optimistic case *model market segmentation*, and in Section 2.4 will show that reasoning about the profit in a world of perfect prediction can provide a useful upper bound on the loss of FBMSP with error.

2.2.1 Key Assumptions and Preliminaries.

As mentioned in the introduction, the optimal feature-based market segmentation and pricing is generally hard to compute. To ensure tractability, in our work we will carry through the common regression assumption that the model error, ϵ, is independent across features i.e. $\mathbb{X} \perp\!\!\!\perp \epsilon$. We consider this assumption to be quite mild, as it underlies many predictive models used in practice, including for example, the well known logit model where a customer's valuation is a linear combination of that customer's features, the offered price, and an idiosyncratic error following a logistic distribution which is *independent* of $\mathbb{X}$. Similar remarks hold for other regression-based models with independent errors. The upshot will be that this necessary assumption for regression is also quite harmonious with pricing, and gives considerable structure and control for analyzing pricing models.

In the next section we will delve into the structure of optimal FBMSP, but first we will illustrate how the independence assumption smooths our problem by considering some related objectives. To this end, consider three auxiliary functions that will later be helpful in analysis of FBMSP, and also are of independent interest.

$$\textbf{Price: } p_\epsilon(x) := \inf\{\arg\max_p p\overline{\mathrm{F}}_\epsilon(p-x)\}, \quad \textbf{Margin: } \theta_\epsilon(x) := p_\epsilon(x) - x,$$

$$\textbf{Revenue: } \mathcal{R}_\epsilon(x) := \max_p p \Pr(x + \epsilon \geq p) = p_\epsilon(x)\overline{\mathrm{F}}_\epsilon(\theta_\epsilon(x)).$$

The price, margin, and revenue functions all serve to model a seller pricing a good for a customer, after having predicted their valuation as $x \in \mathbb{R}$, up to some stochastic error

ϵ. $p_\epsilon(x)$ is the optimal price to offer a customer with predicted valuation x, $\mathcal{R}_\epsilon(x)$ is the revenue of the optimal monopoly price when the valuation distribution is $x + \epsilon$, and $\theta_\epsilon(x)$ is the difference or *margin* between the predicted valuation x and the offered price $p_\epsilon(x)$. Note, $p_\epsilon(x)$ is uniquely defined to be the minimum price that achieves the maximum revenue, such a minimum is necessary since for some distributions ϵ, there may be many prices that maximize the revenue (see Example 1, for an extensive discussion on when the optimal price is unique, or equivalently when the revenue function is strictly unimodal, see [141]).

In the following lemma we summarize some of the structure we observe in these functions.

Lemma 1 (General Properties of $p_\epsilon(\cdot), \theta_\epsilon(\cdot), \mathcal{R}_\epsilon(\cdot)$)**.** *For any distribution ϵ such that $\mathbb{E}[\epsilon] = 0$, the following properties hold:*

(a) $\theta_\epsilon(x)$ *is a decreasing function.*

(b) *For any $0 < x_1 < x_2$, we have*

$$\overline{F}_\epsilon\left(\theta_\epsilon(x_1)\right)(x_2 - x_1) \leq \mathcal{R}_\epsilon(x_2) - \mathcal{R}_\epsilon(x_1) \leq \overline{F}_\epsilon\left(\theta_\epsilon(x_2)\right)(x_2 - x_1).$$

Moreover, for all x such that $p_\epsilon(x)$ is continuous (i.e. $p_\epsilon(x^-) = p_\epsilon(x^+)$), the derivative of $\mathcal{R}_\epsilon(x)$ exists and $\frac{d}{dx}\mathcal{R}_\epsilon(x) = \overline{F}_\epsilon(\theta_\epsilon(x))$.

(c) $\mathcal{R}_\epsilon(x)$ *is increasing, continuous, and convex.*

Lemma 1 implies that independence between the error and valuation model induces prices that result in a monotone increasing sales probabilities via $\theta_\epsilon(x)$, and a convex revenue function $\mathcal{R}_\epsilon(x)$ with interpretable, bounded derivatives. All three parts of the lemma are proved by examining the induced optimal prices, $p_\epsilon(x)$, and noting that $p_\epsilon(x)$ cannot increase very quickly (i.e. super linearly). Unfortunately, our control is not perfect as $p_\epsilon(x)$ can otherwise be quite poorly behaved; there can be many optimal prices for a given valuation, and worse $p_\epsilon(x)$ can be discontinuous at arbitrarily many points x (see Example 1 for an example). The jump discontinuities in $p_\epsilon(x)$ translate directly to non-differentiable points in the revenue function. As we will see in Section 2.3, the structure provided by independence is not quite enough to enable the efficient computation of the optimal FBMSP, however it will be critical in our analysis of such policies.

2.3 Computing Optimal Feature-Based Market Segmentation and Pricing

In this section we study the problem of finding the jointly optimal FBMSP, culminating with conditions and an algorithm under which the optimal policy can be computed efficiently. We will first show that when valuations are drawn from a regression model with general independent residuals, the optimal policy is NP-Hard to compute. We then identify that the hardness stems from some pathological segmentation properties, and define a natural property to characterize nice segmentations which we call *interval*. Our main positive result of this section is to show that when the residuals are log-concave the optimal segmentations are interval, and further, the optimal interval segmentations can be found in cubic time via dynamic programming. Thus, for realistic valuation models under standard regression assumptions, the jointly optimal segmentation and pricing can be directly computed instead of having to resort to heuristic segment-then-price approaches.

2.3.1 Hardness of FBMSP

To understand the computational difficulty of FBMSP, in this subsection we will review a folklore hardness result and show that even under our restricted model of regression based valuations, the problem remains intractable. Our proof of this hardness result will yield guiding intuition for how an additional condition of log-concavity on the residuals should be computationally useful.

First, let us review how the hardness of general market segmentation and pricing problems are typically proved by describing a reduction from k-hitting set[1]. Now, consider the decision version of the hitting set problem given a universe of elements $\mathcal{U} = \{p_1, p_2, \ldots, p_m\}$, a collection of subsets $\mathcal{S} = \{S_1, S_2, \ldots, S_n\}, S_i \subset \mathcal{U}$, $\cup S_i = \mathcal{U}$, and some target $k \in \mathbb{N}$. The objective is to find a set $H \subset \mathcal{U}, |H| \leq k$, that intersects with every subset in the collection i.e. $H \cap S_i \neq \emptyset$ for all $i \in [n]$, or else return that no such set of size k exists. To encode hitting set using a general version of k-market segmentation and pricing, associate with each element p_j, a distinct positive number (for instance, encode each p_i as the number i) and

[1]This correspondence is implicit in the work of [82].

consider an instance with n customers. For each customer, let their valuation be described as a draw from a personal, independent distribution F_i, $i \in [n]$ so that the revenue gained from customer i when they receive price p is $p\overline{\mathrm{F}}_i(p)$. Construct F_i so that the value associated with p_j maximizes $p\overline{\mathrm{F}}_i(p)$ if and only if $p_j \in S_i$. This is easily achieved by appropriately discretizing the *equal revenue* distribution, which is the distribution such that $F(x) = 1 - 1/x$ for $x \in [1, \infty)$. For example, if $S_i = \{1, 3, 4\}$ then a distribution F_i such that $\overline{\mathrm{F}}_i(x) = 1$ for $x \leq 1$, $\overline{\mathrm{F}}_i(x) = 1/3$ for $x \in (1, 3]$, $\overline{\mathrm{F}}_i(x) = \frac{1}{4}$ for $x \in (3, 4]$, and $\overline{\mathrm{F}}_i = 0$ otherwise, is a construction of the desired form, achieving a maximum revenue of 1 only at prices 1, 3 and 4.

To complete the reduction, note the market segmentation and pricing problem faced by the seller is to find a partition of customers into k segments that maximizes revenue. The absolute upper bound on this revenue is $\sum_i \max_p p\overline{\mathrm{F}}_i(p)$ which would be the revenue of personalized pricing. If the seller can achieve this bound with k segments, it must be by using a price in each segment which is optimal for every customer in that segment. This implies that on each segment the intersection of each customer's optimal price set (which are in correspondence with S_i) is non-empty. Taking those optimal prices from each segment thus exactly corresponds to a size k hitting set, and if the optimal revenue from k segments is less than the revenue from personalized pricing then no size k hitting set can exist which completes the argument.

Note, the above archetypal reduction assumed independent valuation distributions for each customer, F_i. That is, to encode k-hitting set required n distinct, independent distributions - one for each customer. This is different than in our model, where all valuations are described by a common regression function with independent but identical noise. Thus in our model all customer valuations are described by just two distributions, one over the features and one over the error. As our model is simpler and less expressive, it is no longer clear if a similar hardness result as above is possible. As we show in the following theorem, the problem remains NP-hard in our model, although the proof requires a significantly more intricate construction.

Theorem 1 (Hardness of FBMSP)**.** *Suppose that $V = \mu(\mathbb{X}) + \epsilon$, $\mathbb{X} \perp\!\!\!\perp \epsilon$, and $\mathbb{E}[\epsilon] = 0$. Then finding the optimal FBMSP policy is NP-hard.*

Sketch of Proof of Theorem 1. The proof of Theorem 1 follows by reduction to hitting set, as in the general case. For every instance of the hitting set problem, we show that there exist estimations of customers' valuation, and prediction error's, such that deciding if there are k (or less) elements that hit all the subsets is equivalent to deciding whether there are k segments and prices such that the total revenue is $\frac{n(n+1)}{2}$, where n is the number of subsets in the hitting set problem (equivalently, the number of customers in the market). Our construction follows by designing a error distribution ϵ which results in a number claw like functions for each customer, that are then spread by translation to encode a set of optimal prices for each valuation level $\mu(\mathbf{x}_i)$. Section 2.3.1 gives an example of our hardness reduction for a small instance. $\square$

Theorem 1 implies it is impossible to solve general FBMSP efficiently if P $\neq$ NP. This leaves us two options, either to look for approximate solutions for general FBMSP, or to enrich the structure of our model by imposing additional assumptions. We briefly explore the former in Remark 1, but will focus mainly on the later.

Remark 1. *While an optimal policy for general error distribution cannot be found in polynomial time, we note that a constant factor approximation to the optimal feature-based market segmentation and pricing is obtainable. Specifically, a* $(1 - 1/e)$ *factor approximate policy can be found in polynomial time since the objective function is positive valued, monotone, and submodular. We formalize this observation in Appendix A.4 in the appendix.* $\square$

Figure 3: Example of hitting set and FBMSP.

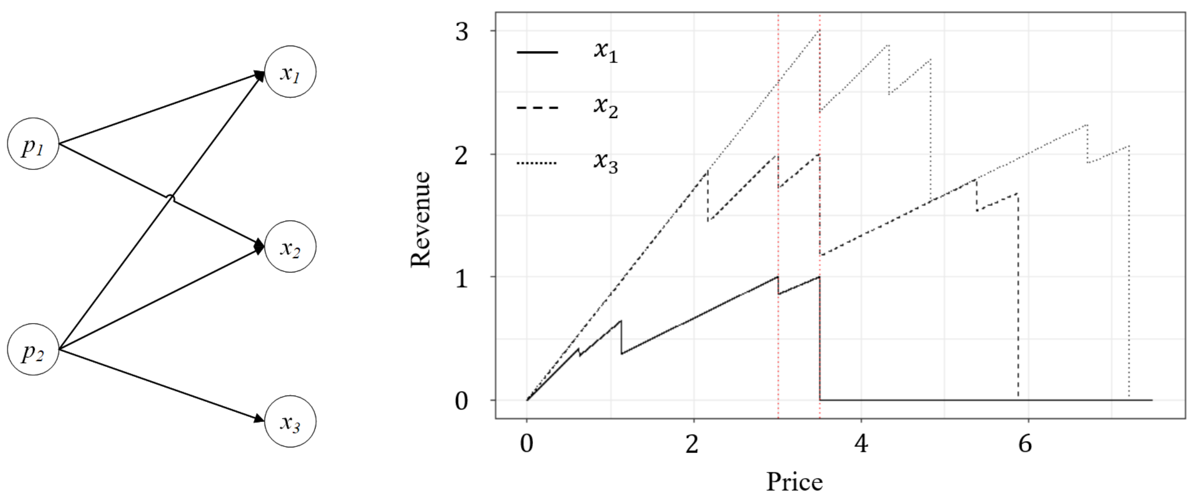

Note. In the left panel is a graph representation of a small instance of a hitting set problem. To illustrate the translation of a hitting set problem to an FBMSP problem, assume we have 3 customers with predicted valuation $x_1 = 3$, $x_2 = \frac{43}{8}$, $x_3 = \frac{161}{24}$. Further, let the estimation error ϵ be supported on $\left\{-\frac{161}{24}, -\frac{77}{5}, -\frac{19}{8}, -\frac{15}{8}, 0, \frac{1}{2}\right\}$ with probability masses $\left\{\frac{1}{7}, \frac{4}{21}, \frac{2}{21}, \frac{5}{21}, \frac{1}{21}, \frac{2}{7}\right\}$. Finally, let $p_1 = 3$ and $p_2 = 3.5$. In the right panel we plot the revenue curves for each valuation x_i. We can see that the revenue for each customer is maximized only at either p_1 or p_2 (red dashed lines) which represent the connections between price nodes and valuation nodes in left panel, i.e., revenue from customer with valuation x_3 is maximized at price $p_2 = 3.5$, revenues from customers with valuations x_2 and x_1 are both maximized at $p_1 = 3$ and $p_2 = 3.5$.

2.3.2 Feature-Based Market Segmentation and Pricing with Log-Concave Residuals

In the previous subsection we studied FBMSP in a setting where the underlying error in the valuation model was arbitrary, and in the proof of Theorem 1 we leveraged this freedom to construct a general error distribution such that it induced jagged, delicately overlapped revenue functions that made the problem intractable. In this section we will consider assumptions that evade such pathological constructions. To that end, recall in Lemma 1 were able to characterize many things about the revenue function $\mathcal{R}_\epsilon(x)$, but less about the structure of the pricing function which we noted could vary (drop) dramatically between similar valuations. It is precisely these discontinuities in $p_\epsilon(x)$ that enable our construction in Section 2.3.1, and give it it's discrete quality that makes it difficult to optimize. A natural question then to ask is, for suitably smooth error distributions/revenue functions, is it still hard to compute the optimal FBMSP?

We will make one additional assumption about the error distribution that enforces such a notion of smoothness. Namely, we will assume the distribution is *log-concave*, a canonical assumption in the pricing and revenue management literature. Note, many standard distributions are log-concave including normal, exponential, uniform distributions, etc.

Definition 1 (Log-Concave Error)**.** *A random variable ϵ with density f_ϵ is* ***log-concave*** *if* $\log(f_\epsilon(x))$ *is a concave function.*

To understand how log-concavity in the error function translates into tractability for FBMSP, we will first show that it implies a continuous, increasing price function $p_\epsilon(x)$, precluding behaviour like in Section 2.3.1. Leveraging this continuity in the prices, we can then show the optimal segmentation must be well structured in the sense that the segmentation groups together customers with similar predicted valuations for the good. Such segmentations are natural, easy to interpret as low/medium/high/etc. type segments, and as we will show, easy to optimize and analyze. We will call segmentations that group together customers with similar valuations *interval*, and define them as follows.

Definition 2 (Interval Segmentation)**.** *We will call a segmentation,* $\{\mathcal{X}_i\}_1^k$*, an* ***interval*** *segmentation if there exists real numbers* $0 < s_0 \leq s_1 \leq \ldots \leq s_k = \sup_{\mathbf{x}} \mu(\mathbf{x})$ *such that each segment* $\mathcal{X}_i$ *can be written as* $\mathcal{X}_i = \{\mathbf{x} | \mu(\mathbf{x}) \in [s_{i-1}, s_i)\}$.

When describing interval segmentations, we will often denote the segmentation by just the end points of the intervals in the valuation space that define them, $\{s_i\}_0^k$. We emphasize that not all optimal market segmentations are interval, certainly the ones induced by the construction in Theorem 1 are not, but also even simple error distributions can have more complicated structure as we demonstrate in Example 2. Thankfully, it turns out the smooth notion of error captured by log-concavity, and the intuitive structure of interval segmentations are harmonious notions. In the following lemma we show that log-concavity in the error removes any jump discontinuities from the price function, which in turn allows us to prove that the revenue optimal FBMSP is interval.

Lemma 2 (Properties of Log-Concave Error)**.** *Suppose that $V = \mu(\mathbb{X}) + \epsilon$ where ϵ is log-concave, $\mathbb{X} \perp\!\!\!\perp \epsilon$, and $\mathbb{E}[\epsilon] = 0$. Then,*

(a) *$p_\epsilon(x)$ is an increasing and continuous function.*

(b) *The optimal segmentation is interval.*

(c) *The price on each segment $p_\epsilon(\mathcal{X}_i)$ equals $p_\epsilon(\mu(\mathbf{x}))$ for some $\mathbf{x}$ such that $\min_{\mathbf{x}\in\mathcal{X}_i} \mu(\mathbf{x}) \leq \mu(\mathbf{x}) \leq \max_{\mathbf{x}\in\mathcal{X}_i} \mu(\mathbf{x})$.*

Lemma 2 shows that, by assuming the error in the regression model is log-concave, all the previously mentioned pathologies vanish. First, we show that $p_\epsilon(x)$ becomes a strictly increasing function which, combined with Lemma 1, implies that the revenue function $\mathcal{R}_\epsilon(x)$ is differentiable everywhere, and it's derivative is simply the sale probability. We then show in (b) that the upshot of this additional smoothness for FBMSP is that the segmentation policy becomes interval. Moreover, the optimal price to offer on each segment is contained *in the segment*, as the optimal price for some feature vector. This locality of the price and segment then enables fast computation of the optimal policy via dynamic programming, as we describe next in the main theorem for this section.

Theorem 2 (Computing Feature-Based Market Segmentation)**.** *Suppose that $V = \mu(\mathbb{X}) + \epsilon$ where ϵ is log-concave, $\mathbb{X} \perp\!\!\!\perp \epsilon$, and $\mathbb{E}[\epsilon] = 0$. Let $n = |supp(\mu(\mathbb{X}))|$, and suppose $p_\epsilon(\mathcal{X}_i)$ for any fixed segment $\mathcal{X}_i$ can be computed in time m_ϵ. Then the optimal feature-based market segmentation can be computed in $O(n^2(k + m_\epsilon))$.*

Theorem 2 is our main result, and states that by leveraging the structural properties in Lemma 2, the optimal policy can be computed quickly and efficiently in terms of the size of the support of the regression model. Note, we assume $\mu(\mathbb{X})$ is finitely supported, and believe this is natural and corresponds to simply running the regression model back over the sample of customer outcomes which were used to generate the model. Further, we assume the running time to compute $p_\epsilon(\mathcal{X}_i)$ as a subroutine is bounded by some number that is related only to ϵ. Again, we believe this assumption is natural since when ϵ is log-concave, the revenue function for some sample $p\Pr(x+\epsilon \geq p)$ is unimodal in p, and the price of the segment $p_\epsilon(\mathcal{X}_i)$ can be computed simply running a binary search for the optimal price on the range of prices $[p_\epsilon(\min_{\mathbf{x}\in\mathcal{X}_i} \mu(\mathbf{x})), p_\epsilon(\max_{\mathbf{x}\in\mathcal{X}_i} \mu(\mathbf{x}))]$ by Lemma 2(c). Lastly, we note that when the error distribution is discrete, there is a corresponding notion of *discrete log-concavity* [112], under which our results continue hold without modification.

In this section, we have characterized when and how we can compute FBMSP optimally, in the subsequent sections we turn our attention to tuning and implementing it as a revenue management strategy.

2.4 Analyzing Feature-Based Market Segmentation and Pricing

In the previous section, we studied how to compute the optimal FBMSP under some assumptions, for a given number of segments/prices k. We showed that while, in general, it is hard to do so, in the important and realistic case when error is log-concave, the optimal policy has an intuitive structure that allows for easy computation. In this section, we continue to build on the structural insights of the last section, and show that beyond just computation, optimal FBMSP inherits a number of attractive properties and performance guarantees that may help guide practitioners in implementing such policies, and particularly in deciding how many segments to use.

Throughout this section, when the underlying valuation model varies we will use a superscript to explicitly identify the valuation distribution with which the revenue is computed i.e., $\mathcal{R}_{kXP}^{\mu(\mathbb{X})+\epsilon}$ is the standard revenue of FBMSP, $\mathcal{R}_{kXP}^{\mu(\mathbb{X})}$ is the revenue of FBMSP with no error in valuation model, and so on.

2.4.1 FBMSP vs. Feature-Based Personalized Pricing

In this subsection, we study the relative gaps between the optimal FBMSP and the natural upper bound of feature-based personalized pricing, paying close attention to how this gap informs a good choice of k. As mentioned in the introduction, FBMSP closely resembles real-world data-driven semi-personalized pricing strategies where sellers are constrained in the number of the segments/prices they can offer. Specifically, in FBMSP the number of prices and segments is capped at k, whereas feature-based personalized pricing is equivalent to FBMSP when $k \to \infty$. In fact, for any market where valuations are distributed according $V = \mu(\mathbb{X}) + \epsilon$, and $\mathbb{X} \perp\!\!\!\perp \epsilon$, the revenue of a seller implementing feature-based personalized pricing can be succinctly described as an expectation over the revenue function i.e., $\lim_{k\to\infty} \mathcal{R}_{kXP}^{\mu(\mathbb{X})+\epsilon} := \mathbb{E}_{\mathbb{X}\sim F_{\mathbb{X}}}[\mathcal{R}_\epsilon(\mu(\mathbb{X}))]$, since the seller offers the optimal price for each context $\mathbf{x}$ which garners revenue $\mathcal{R}_\epsilon(\mu(\mathbf{x}))$.

Intuitively then, a good choice of k should be one that is not too large, so as to be implementable, but one that is still close to the maximum achieveable revenue of personalized pricing, i.e., one that shrinks the gap,

$$\mathbb{E}_{\mathbb{X}\sim F_{\mathbb{X}}}[\mathcal{R}_\epsilon(\mu(\mathbb{X}))] - \mathcal{R}_{kXP}^{\mu(\mathbb{X})+\epsilon}. \tag{2}$$

One difficulty that may be encountered when attempting to choose k to reduce Eq. (2) is that it is sensitive to the error distribution ϵ, which may be hard to know precisely or require extensive market research to obtain. It may be preferable for an analyst attempting to choose k to work with an upper bound on this difference that is agnostic to the true error distribution. Interestingly, by assuming there is no error an analyst can achieve precisely such an upper bound. In the following theorem we show that one can bound the loss between FBMSP and feature-based personalized pricing by examining the gap between the two policies when ϵ is assumed to be 0 a.s. We will refer to this loss as *model market loss*, since it depends only on $\mu(\mathbb{X})$ and not the underlying error distribution.

Theorem 3 (Model Loss vs. True Loss)**.** *Suppose $V = \mu(\mathbb{X}) + \epsilon$, $\mathbb{X} \perp\!\!\!\perp \epsilon$, and $\mathbb{E}[\epsilon] = 0$. Then,*

$$\underbrace{\mathbb{E}_{\mathbb{X}\sim F_{\mathbb{X}}}[\mathcal{R}_\epsilon(\mu(\mathbb{X}))] - \mathcal{R}_{kXP}^{\mu(\mathbb{X})+\epsilon}}_{\textit{Actual Market Loss}} \leq \underbrace{\mathbb{E}_{X\sim F_{\mathbb{X}}}[\mu(X)] - \mathcal{R}_{kXP}^{\mu(\mathbb{X})}}_{\textit{Model Market Loss}} .$$

Theorem 3 gives a theoretical foundation through which an analyst can analyze the performance of FBMSP for various k directly in the model without worrying about the particular form of the error distribution. Whatever loss is perceived in the model market bounds the true loss in practice automatically.

Further, we note that the proof of Theorem 3 is constructive, and implies a simple heuristic for setting feature-based market segmentation and pricing strategies when ϵ is not log-concave, or ϵ is unknown. In these instances, a seller can simply compute the optimal k-FBMSP letting ϵ be 0 a.s. In this situation, the optimal policy is interval and can be described by segmentation end points $\{s_i\}_{i=0}^k$ on the model market $\mu(\mathbb{X})$, which can be used to generate the segments $\mathcal{X}_i$. From those segments, since ϵ is either unknown or not tractable to work with computationally, the firm can instead perform price experimentation to learn the prices that maximize $p_\epsilon(\mathcal{X}_i)\Pr\left(s_i+\epsilon \geq p_\epsilon(\mathcal{X}_i)\right)$, and offer that price on each segment. While both the partition into segments $\{\mathcal{X}_i\}_{i=1}^k$, and the prices offered on each segment $\{p_\epsilon(\mathcal{X}_i)\}_{i=1}^k$ may be sub-optimal under the true error, such a strategy is guaranteed to earn more than $\mathcal{R}_{kXP}^{\mu(\mathbb{X})+\epsilon} + \mathcal{R}_{kxP}^{\mu(\mathbb{X})} - \mathbb{E}_{X\sim F_{\mathbb{X}}}[\mu(X)]$ by rearranging Theorem 3, and this guarantee smoothly tends to the optimum as the error in the model diminishes.

Theorem 3 allows an analyst to search for a choice of k without referring to the error distribution, a next natural question to ask is then, how long can this search take? That is, at what rate does $\mathcal{R}_{kXP}$ to converge to the revenue of feature-based personalized pricing? In our next theorem we show this convergence is linear in k, and quite fast when the range of valuations is not too wide.

Theorem 4 (Bounded Loss with k Segments)**.** *Suppose* $V=\mu(\mathbb{X})+\epsilon$, $\mathbb{X} \perp\!\!\!\perp \epsilon$, *and* $\mathbb{E}[\epsilon]=0$. *Let* $L=\inf_{\mathbf{x}}\mu\left(\mathbf{x}\right)$ *and* $U=\sup_{\mathbf{x}}\mu\left(\mathbf{x}\right)$, *then*

$$\mathbb{E}_{\mathbb{X}\sim F_{\mathbb{X}}}[\mathcal{R}_\epsilon(\mu\left(\mathbb{X}\right))] - \mathcal{R}_{kXP}^{\mu(\mathbb{X})+\epsilon} \leq \frac{U-L}{k}.$$

The proof of Theorem 4 constructs a (suboptimal) segmentation strategy by equally partitioning the quantile space. Interestingly, the dependence $O\left(\frac{1}{k}\right)$ appears typical for many valuation distributions, as we plot in Fig. 4. Intuitively, this behavior can be explained in the following way. As we segment into smaller pieces, any distribution with a smooth density appears locally uniform on each segment. Example 3 establishes that the convergence rate for a uniform matches Theorem 4 up to constant factors, suggesting that, at least for large k, the rate should also be approximately tight for many distributions.

Figure 4: Difference between FBMSP and feature-based personalized pricing for standard distributions.

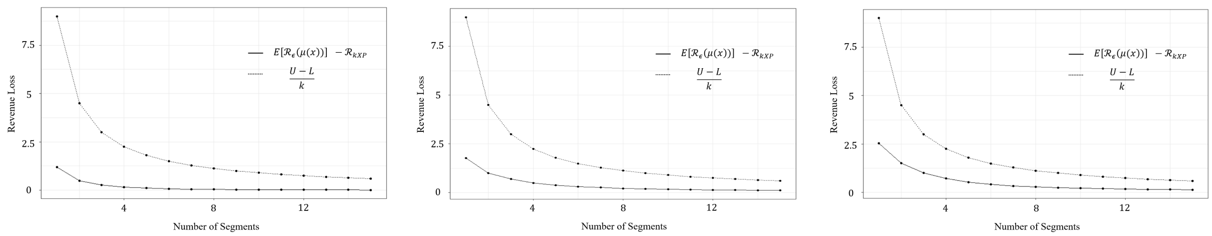

Note. Depicted is the revenue loss versus the number of segments for three standard error distributions, and when predicted customer valuations are drawn uniformly from $[1, 10]$ i.e., $\mu(\mathbb{X}) \sim \text{Uniform}[1, 10]$. For simplicity, in the plots we use $\mathcal{R}_{kXP}$ and $\mathbb{E}[\mathcal{R}_\epsilon(\mu(\mathbb{X}))]$ to denote $\mathcal{R}_{kXP}^{\mu(\mathbb{X})+\epsilon}$ and $\mathbb{E}_{\mathbb{X}\sim F_{\mathbb{X}}}[\mathcal{R}_\epsilon(\mu(\mathbb{X}))]$, respectively. In the left panel, the prediction error $\epsilon \sim \mathcal{N}(0, 1)$. In the middle panel, the prediction error $\epsilon \sim \text{Uniform}[-1, 1]$. In the right panel, the prediction error follows a Weibull distribution where the shape parameter is 5 and scale parameter is 1. In each panel, we plot the revenue loss versus the bound $\frac{U-L}{k}$ in Theorem 4. We note the convergence rate for all three error distributions appears to be $\Theta(1/k)$.

2.4.2 Revenue Concavity in the Number of Segments

Theorems 3 and 4 give an analyst insight into how to handle the error when searching for k, and a bound on how large a k may be needed to achieve a desired level of revenue loss. In the final result of this section, we show a nice structural property of the optimal revenue that an analyst can use to further hone their search for k. Specifically, in Theorem 5 we show that when the residual is log-concave, the revenue of FBMSP is concave in the number of segments.

Theorem 5 (Segmentation Concavity)**.** *Suppose that $V = \mu(\mathbb{X}) + \epsilon$ where ϵ is log-concave, $\mathbb{X} \perp\!\!\!\perp \epsilon$, and $\mathbb{E}[\epsilon] = 0$. Then $\{\mathcal{R}_{(k+1)XP} - \mathcal{R}_{kXP}\}_{k=1}^{\infty}$ is a non-increasing sequence.*

Concavity in terms of k has the operational interpretation that the revenue garnered by additional segments has diminishing marginal returns. Such a property is not guaranteed in general, and especially not for heuristic segment-then-price approaches, as we will see in Section 2.5. Moreover, combining Theorems 3 and 4 and Theorem 5 together allows an analyst to search for the best choice of k via an *elbow method* [23] on the model market. Such a method results in a k with a provable guarantee on the loss (Theorem 3), is likely quite small (Theorem 4), and further the elbow will be unique (Theorem 5).

In the next section we implement our segmentation and pricing strategy on real data and illustrate it's advantages versus segment-then-price heuristics.

2.5 Case Study: Setting Mortgage Interest Rates

In Section 2.3, we showed how to find jointly optimal FBMSP when a seller has trained a regression-based valuation model with independent, log-concave residuals. Then in Section 2.4, we provided a set of results to aid in the analysis of FBMSP policies and guide the choice of k, the number of segments/prices. In this section, we perform a case study to highlight some features of our approach, which we compare and contrast with prominent heuristic approaches for segment-then-price (STP). Specifically, using a real data set of home mortgage offers and acceptances in Pennsylvania in 2020, we build a probit regression model to predict the probability that an applicant will take a mortgage at an offered interest rate. Next, we transform the probit regression model into a model of customer valuation measured as the maximum interest rate they will accept. We then compare our optimal method for FBMSP with STP via a number of different simulations on the data set. All data and code for this section are publicly available at [**Blinded for Review**].

2.5.1 Description of Data Set

Our case study is based on a dataset collected in accordance with the Home Mortgage Disclosure Act (HMDA) (the HMDA website where the data is hosted is `https://ffiec.cfpb.gov/`). Specifically, we downloaded the data provided by all financial institutions

Variable	Type	Description and Statistics
Action taken	Binary	The action taken on the covered loan or application • 1 (accepted), Frequency = 11491, Percent = 77.0% • 0 (rejected), Frequency = 3425, Percent = 23.0%
Interest rate	Continuous	The interest rate for the covered loan or application (%) • Mean = 3.4%, Std = 0.9%
Income	Continuous	Applicant's gross annual income (in thousands of dollars) • Mean = 110.08, Std = 94.76
Marital status	Binary	Marital status derived from applicants' marital status fields • 1 (joint), Frequency = 5799, Percent = 38.9% • 0 (male or female), Frequency = 9117, Percent = 61.1%

Table 1: Descriptions and summary statistics for explanatory variables in our home mortgage dataset.

in Pennsylvania who had offered a loan for the purpose of enabling a home purchase in 2020. The dataset consists of information about the applications, including demographic information about the applicant, their income level, the loan amount the bank offered, the interest rate, and whether or not it was accepted. After removing unsuitable rows (rows with data missing or extreme outliers), there were 14,916 approved applications in total, and 11,491 (77%) of the approved applications resulted in a loan that was accepted at the bank offered interest rate. Table 1 summarizes the variables (features of customers) we use in our case study.

As a preliminary, note we can think of the interest rate the bank offers on the loan as a take-it-or-leave-it price, and the customers choice whether or not to accept the loan as a decision to purchase or not purchase at a given price. Using the price variation in this data, we will estimate a customer valuation model so that we can train and evaluate market segmentation and pricing models. In the first step, we use a probit regression model to predict the probability that a customer will take the offered interest rate. Table 2 shows the coefficient estimates for the probit regression model, and Fig. 21 shows the prediction of the probability a customer will take the approved application for a given interest rate. We then transform our probit model into a linear valuation model of the form $V|\mathbf{x} = \mu(\mathbf{x}) + \mathcal{N}(0, \sigma)$, where $\mathbf{x}$ is a feature vector including the interest rate, income level of the customer, and demographic information, and $\mu(\mathbf{x}) = \mathbf{x}^t\beta$. Our transformation from probit regression model to linear valuation model follows [30], for a short primer describing such transformations, see Appendix A.3 in the appendix. All subsequent feature-based segmentation and pricing policies will be based on this derived linear valuation model, $V|\mathbf{x} = \mu(\mathbf{x}) + \mathcal{N}(0, \sigma)$.

2.5.2 Comparison with Segment-then-Price

To assess the real impact of our optimal FBMSP policies, we will compare against heuristic segment-then-price (STP) policies. For STP, we will segment customers using the popular k-mediods algorithm [106, 114, 115] with Gower distance [60, 61]. The price optimization is then done over the found segments, and can be computed in polynomial time for error with finite support.

First, we examine the segments generated by STP compared with those from optimal FBMSP. Fig. 5 shows that STP will group customers first based on differences in gender and/or race. Gender and race are certainly heterogeneous across our data set, however, these differences are not necessarily the distinctions that are revenue-maximizing to delineate on. In comparison, optimal FBMSP will segment customers into different groups based on their valuations, which is only weakly correlated with gender/race in our data. Therefore, compared to STP, FBMSP will not only achieve better revenue but does so in an explainable way by grouping customers with similar predicted valuations, instead of merely similar

Variable	Estimate	Std. Error	z value	$\Pr(>|z|)$	Significance
(Intercept)	3.8869	0.1126	34.5314	$< 2.2 \times 10^{-16}$	***
Interest Rate	-0.8704	0.0277	-31.4735	$< 2.2 \times 10^{-16}$	***
Income	-0.0009	0.0002	-4.1142	3.885×10^{-05}	***
Derived race	0.4709	0.0545	8.6374	$< 2.2 \times 10^{-16}$	***
Derived gender	0.1721	0.0441	3.9003	9.606×10^{-5}	***

Table 2: Probit regression coefficients. Significance levels: ***: < 0.001, **: < 0.01, *:$<$ 0.05.

demographic features which may have negative social or legal ramifications.

Figure 5: Segments for STP and FBMSP.

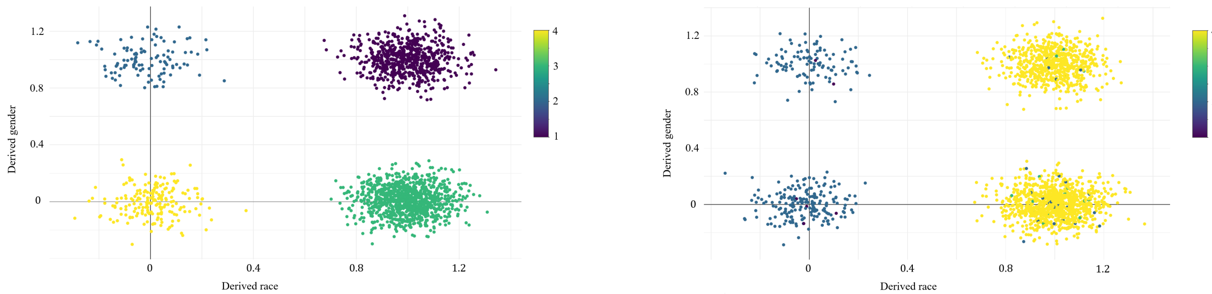

Note. Here we plot segments for STP and FBMSP when the number of segments $k = 4$. Since both derived gender and derived race are binary variables, we add some random noise to each point for clarity of presentation (without noise, all points of the same color would be on top of one another in the left panel). In the left panel, the segments are obtained using k-medoids algorithm. In the right panel, we use dynamic programming to do optimal FBMSP.

Figure 6: Average interest comparison between STP and FBMSP.

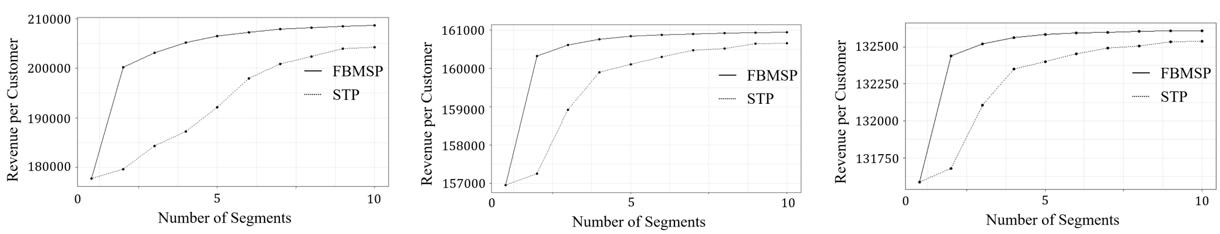

Note. Here we plot the average interest per customer (i.e. expected revenue) for STP and FBMSP, for different levels of prediction error in our valuation model. In the left panel, the standard deviation of prediction error is $\sigma = 0$, in the middle panel, the standard deviation of prediction error is $\sigma = 0.5$, in the right panel, the standard deviation of prediction error is $\sigma = 1$.

To compare the difference in revenue garnered by STP and FBMSP, we will examine the difference in total interest a customer will pay on average over the lifetime of the loan (i.e. average revenue per customer), where the interest is calculated using the standard fixed monthly payment formula ([31]). In Fig. 6, we plot the expected revenue the firm can get from one customer on average (across all segments), against the number of segments. We first note that the revenue per customer is increasing for both FBMSP and STP model. However for FBMSP, the revenue the firm can get from each customer on average is concave in the number of segments, whereas in the STP model this is clearly not the case. In our case study, STP often gets "stuck" at small choices of k, and requires a 3+ of segments before it can achieve strong revenue, whereas FBMSP is guaranteed to get the most revenue out of a small number of segments. We see the differences between FBMSP and STP are most pronounced when only a small number of segments are used, which is precisely the case of interest in industry. We further note that for both models, while smaller error in the prediction model will yield higher revenue, the gap between the two strategies is also more pronounced when the error is small, suggesting that for sophisticated firms with high quality feature data the benefits of FBMSP are even greater.

2.5.3 Finding the Optimal Number of Segments using Regression Model

One additional benefit of the concavity of optimal FBMSP is it enables us to easily choose the number of segments via the elbow method heuristic. The elbow method is the most commonly used heuristic for finding the optimal number of segments for unsupervised learning. The intuition is that one should choose a number of segments so that adding another segment doesn't give a much better modeling of the data (see [23] for more discussion about the elbow method and its applications). To use the elbow method, one prerequisite is that the objective function is monotone in the number of segments. In general, the objective function, revenue per customer, is not necessarily even increasing for the STP. Unlike STP, in Theorem 5 we showed that the revenue is concave in the number of segments k. At some value for k, the revenue increases dramatically, and after that, it reaches a plateau, and increasing the number of segments does not dramatically increase revenue. In Fig. 6, for our FBMSP model, 2 or 3 is the elbow of the revenue per customer vs. k plot, whereas for STP, the possible elbow is 4+, and sometimes as large as 8 or 9 which is a prohibitively large number of segments in practice.

2.6 Conclusions

Our framework for computing and analyzing semi-personalized, feature-based market segmentation and pricing allows sellers to personalize prices for customers at increasingly fine levels. However, implementing such tools comes with a significant investment cost in technology, data scientists, and marketing. Motivated by this trade-off and the desire to improve on common heuristic approaches, we examined market segmentation and pricing under some realistic assumptions about how the seller predicts customer valuations. Namely, we define and study the feature-based market segmentation and pricing problem, where sellers have access to a trained regression model. For the general case first prove that computing the optimal feature-based market segmentation and pricing is NP-hard and provide a $(1 - 1/e)$ approximation algorithm. We then show that with the additional assumption of log-concavity of the prediction error, the optimal policy has a simple interval structure that can be computed in quadratic-time via dynamic programming.

We then analyze the properties of optimal feature-based market segmentation and pricing. We show that the loss of k-FBMSP versus a fully personalized pricing bench market can be upper bounded by the (noiseless) model market loss, and decays at a tight rate of $\Theta(1/k)$. We also showed the revenue of optimal FBMSP is concave in the number of segments k. Taken all together, this analysis enables practitioners to find the most suitable k by a simple elbow method, and without loss of much revenue.

Overall, our work seeks to deepen our understanding of semi-personalized pricing strategies, and demonstrate that they are computable, and effective when compared to complicated fully personalized pricing strategies. There are many interesting and important directions left to consider for future work, we highlight three of them here. First, this chapter assumes the production cost of the good is uniform over all segments. Follow-up work may consider heterogeneous production costs among different segments, and ask whether the optimal FBMSP in this case still uses interval segments when the residuals are log-concave.

Second, it may also be interesting to consider the approximation ratio for *interval* segmentations facing general error distributions. Example 2 demonstrates that interval segmentation are not optimal for general error distributions, but how far it is from the optimal segmentation in the worst case is unknown. We emphasize that the $1 - 1/e$ approximation algorithm presented in Remark 1 does not compute interval segmentations, and it may indeed be the case that the optimal interval segmentation (which can always be computed in polynomial time via Theorem 2) could achieve a stronger approximation guarantee. Finally, we assume the firm can charge customers in different segments any segment level price. In practice, the firm may only be able to offer a price menu for customers to choose from. Future work may consider models similar to FBMSP, where customers react to and choose from a size k price menu.

3.0 Pricing Strategies for Online Dating Platforms

3.1 Introduction

More and more people, especially young people, are meeting and falling in love online. Over the past two decades, online dating has displaced conventional mediums such as family, school, or the workplace, to become the most common way for new couples to meet [117]. As of 2019, three-in-ten U.S. adults say they have used a dating site or app before, and that percentage rises to 48% for 18-to-29 year-olds [128]. Moreover, dating apps have been enormously successful in connecting people otherwise left out of traditional dating culture. A full two-thirds of lesbian, gay, or bisexual Americans report using dating apps [127] and, with the COVID-19 pandemic severely restricting the venues in which people can meet, many leading dating apps are reporting record numbers of users and subscribers [96].

The companies that run these apps have in turn grown with the increased demand. The dating services industry has swelled over the last five years, with an annualized growth rate of 12.9% [72], and revenues projected to rise 9.3% in 2021 to $5.3 billion as mobile services expand. In the United States, dating services are largely consolidated under one corporate entity, the Match Group, which as of 2020 is estimated to have cornered more than 60% of the dating app market with its suite of apps, including *Tinder*, *Hinge*, *OkCupid*, and *Match* [96].

The majority of dating apps in the United States, including all apps owned by the Match Group, share a number of structural similarities. First, they are *swipe apps*, where a user on the dating platform is shown a curated queue of potential matches, and interacts with the site by swiping left (to deny) or right (to accept) through the queue. Only after two users mutually swipe right on each other can they then begin to communicate [27]. Second, they follow a subscription-based profit model where users pay a subscription fee per period (e.g. per week, per month, see Figs. 7 and 24 for examples) to use the full features of the platform.

While subscription-based policies are an extremely common way to monetize digital apps, their use for online dating comes with some unique challenges. For one, the value a user obtains from a dating app is inherently probabilistic, it stems from the opportunities to meet someone the app provides. Consequently, at least in the case of a user looking for a permanent partner, the longer a user is on the app, the longer that user is paying for a service that has not resulted in the outcomes they desire. This is coupled with a requirement that, to effectively use a dating app, the user has to actively search and filter potential candidates which can be both time intensive and emotionally fatiguing. Indeed, emotional fatigue or burnout is reportedly experienced by as many as four out of five online dating users, and is a primary cause of subscription cancellations [104]. Thus, for online dating we expect user valuation for the service to actually decay over time. This differs significantly from more conventional subscription-based services like Netflix or Spotify, where satisfied users are more likely to continue paying for the service in the future.

Moreover, in the subscription model, the online dating platform's interest may be in conflict with its users'. As [134] note, a profit-maximizing subscription-based dating app will strive to retain all of its subscribers. However, upon finding a compatible partner, users will terminate their dating app subscription and leave the platform. Thus, the platform has an incentive not to always provide the best potential matches to its users so as to extend their time using the app. The severity of this incentive incompatibility depends on the length of the subscription period for the platform. For instance, if the period length is a lifetime (i.e. a lifetime membership), such incentive issues vanish as the site can earn no further profit from the user after the initial payment.

In the context of online dating some matchmaking sites such as `selectivesearch.com` and `bumble.com` have already adopted lifetime memberships, what we refer to as a *long period* pricing in this chapter, as an attempt to solve these incentive issues. In the pricing section of the Selective Service site (`https://www.selectivesearch.com/pricing`) they state: "*While most dating apps and services are incentivized to keep members paying ongoing fees, Selective Search works with clients through a defined contract. Each contract is for a finite number of introductions over a defined period of time*". Ostensibly, sudden departure from the platform, either from burn-out or from matching, is no issue at all if users buy a lifetime pass. In Fig. 7, examples of both subscription priced (at various subscription period lengths) and lifetime long period priced dating apps are shown for reference.

Figure 7: Examples of subscription and long period pricing in online dating platforms.

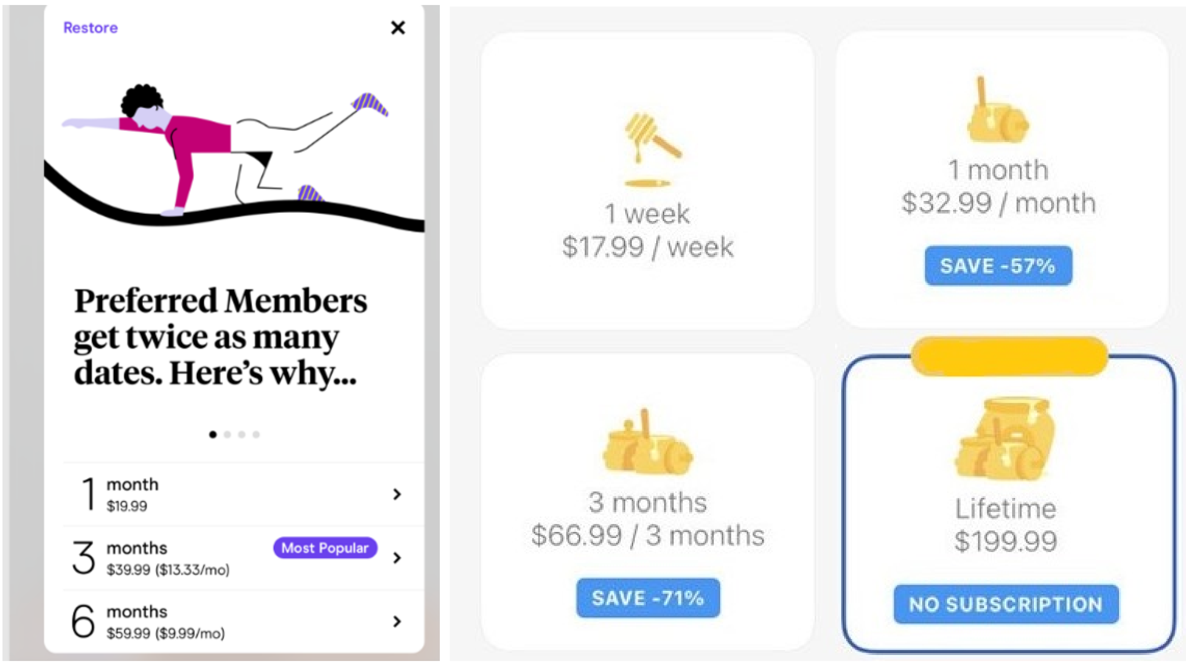

Note. Depicted are pricing options for two popular online dating platforms. On the left is an example of the subscription options at `https://hinge.co/`. On the right is an example of long period pricing (denoted Lifetime) at `https://bumble.com/`, among other options.

In light of the above, in this work we study the monetization of a dating platform (swipe app) run by a profit maximizing monopolist. We focus specifically on the length of the subscription period, and attempt to understand the benefits and trade-offs of shorter or longer periods. We will give special attention to the following two policies which represent the extremes of subscription-based pricing: (i) *short period pricing* (SP), where the online dating platform commits to a fixed price p, and users pay the price in a continuous fashion until they leave the platform, and (ii) *long period pricing* (CP), where the online dating platform **C**ommits to a fixed, one time price p. SP is very flexible and requires no commitment between the platform and user, making it easy to implement, understand, and deploy, even if it leads to possibly confusing incentives. CP is less flexible but may lead to simpler and more harmonious user-platform interaction.

3.1.1 Our Contributions

Given the ostensible advantages and disadvantages of SP and CP for online dating platforms, in this work, we investigate the profit and welfare trade-offs associated with each. A summary of our key contributions and findings is as follows:

1. We give a novel, natural model to describe the operations of an online dating platform. In our model, we study the optimal profit that a platform can obtain from a subscription policy parameterized by a period length, L, and a subscription price, p. We examine two important extremes of this parameterization, short period pricing ($L \to 0$) and long period pricing ($L \to \infty$). We show that short period pricing is robust in the sense that it achieves a constant factor ($\frac{1}{1+e} = 26.8\%$) of the profit of the best-in-class policy for *all* distributions and market parameters, and no other policy can achieve such a guarantee (cf. Theorem 6).
2. Having identified the robust policy, we next take a fine-grained approach and incorporate some additional information about online dating markets that enable us to pin down an optimal policy. Specifically, when the marginal cost of operation is vanishingly small we find that long period pricing achieves both higher profit, and simultaneously matches a higher proportion of the user-base (cf. Theorems 7 and 8). Thus, in well-established online settings with low marginal cost, the use of long period pricing exhibits a "win-win" for both the platform and its users.
3. Finally, we consider the scenario where the platform can incorporate heterogeneous potential match information to vary the match rate. We show that for short period pricing, the platform is incentivized to offer its users the worst possible potential matches first (cf. Theorem 9(a)). On the other hand, under long period pricing the incentive is for the platform to match the user as quickly as possible, and these cost savings are passed on to the user via a lower price (cf. Theorem 9(b,c).

Taken together, our work gives practical insights for market designers looking to improve current swipe based online dating platforms. For new or up-and-coming dating platforms we characterize the robust policy, short period pricing, which can be thought of as an endorsement for new apps to roll out a *freemium* version where fees collected in the form of advertising, and paid in an approximately continuous fashion via user attention. However, for more established apps where the marginal cost of operating with one more user is practically nil, we show a switch to longer period pricing can not only improve revenue, but also increase user welfare via a higher match rate, as well as align the incentives of the user and the platform for harmonious future operations.

3.1.2 Literature Review

Our work is related to several streams of literature in economics, computer science, and operations management. Here, we overview some of these streams and connect them to our work.

3.1.2.1 Analysis of Two-Sided Markets

Online dating is a traditional application area for the theory of two-sided markets. Recently, [11] utilized a game-theoretic model to investigate how platform design can reduce the effort required to find a suitable partner in a two-sided market. Similarly, [78] studied the optimal design of matching platforms by introducing a stylized dynamic fluid model for the two-sided matching with strategic agents. Related two-sided market design questions have also been recently studied by [6, 73, 16, 76, 107] in models with strategic agents. In contrast to these works, in this chapter we treat the market as essentially one sided. We believe this modeling choice better reflects swipe apps as they are currently, since the important two-sided aspects of matching markets are suppressed in our context by the fact that users in swipe apps cannot search over potential matches. Users are shown whoever the platform chooses to show them (or no one at all). In this work, we fully abstract away the mechanics of coordinating matches, and instead analyze directly the pricing strategies of a profit maximizing matching platform.

3.1.2.2 Platform Design in Operations Management and Marketing

Our work contributes to a deep literature dealing with aspects of platform design using models from operations management. The most thematically relevant paper for our work is [135], who study competing matchmakers in a two-period, two-user model with Hotelling valuations. They model dating platforms as strategically investing in matching technologies, investigating the interaction between competition, and providing the best service for their users. [135] does not resolve the suitor-matchmaker incentive issue, but does argue that competition and perfect information about match quality can induce matchmaking platforms to act in the agents' best interests. [48, 49, 46, 64, 17] also study online platform incentives to provide less-than-perfect services. As the US dating market is largely non-competitive [58], we instead focus on changing the pricing structure itself to understand profit, welfare, and incentive trade-offs, and in a significantly more general model.

3.1.2.3 Pricing Strategy in Operations Management

The closest to this chapter, in terms of the framework and style of pricing strategy analysis, is [85]. They consider two business models, pay-per-use selling and product selling, which roughly correspond to our short period pricing and long period pricing, respectively. They focus primarily on an equilibrium analysis of the business model choice under duopoly. Their work explores the scenarios in which the pay-per-use model is more profitable than product selling. Other analyses of pay-per-use selling and product selling can be seen in [133, 124, 29, 4, 15]. Analysis of similar business models can be seen in [101]. In our work, we analyze the pricing strategies of profit maximizing monopolists running an online dating app and model the unique aspects of this emerging application.

3.1.2.4 Other Matchmaking Markets

Outside of dating/marriage, online matchmaking is also an increasingly well-studied problem for labor markets [25, 19], which is the next closest application to our work. Online matchmaking has also been studied in the video game industry. [34] study the problem of maximizing player engagement in video games through improved matchmaking. They focus on a stylized model with different skill levels of players, and where winning or losing influences the players' willingness to stay on the platform. [37, 71, 43] also investigate how to improve players' engagement through matching.

3.1.3 Chapter Outline

In Section 3.2, we introduce our model and describe the profit achievable from a pricing strategy parameterized by a period length L, and a subscription price, p. In Section 3.3, we prove a best-in-class profit guarantee for SP. In Section 3.4, we consider the case when marginal costs are vanishingly small and show CP is both optimal with respect to profit, and yields higher match proportions. In Section 3.5, we extend our model to the case when the match rate can vary. Finally, in Section 3.6, we discuss the implications of our work for online dating platforms, their users, and potential regulators, as well as highlight some interesting avenues for future research.

3.2 Model and Preliminaries

We consider a profit maximizing online dating platform serving a market of users seeking permanent partnership, which we refer to as a *match*. While colloquially a "match" may refer to someone who has only shown preliminary interest in the user, in this work we will exclusively use *match* or *matching* to mean the successful formation of a long-term partnership culminating in a departure from the platform. We model a random user's valuation (willingness-to-pay) for matching as a non-negative random variable V drawn from a known distribution F with density f. The user's valuation for matching decays over time at a constant rate $\delta \in (0, 1)$, so the user's valuation/willingness-to-pay for matching after t time on the platform is $V_t := \delta^t V$. We use the notation $\overline{\mathrm{F}}(x) := 1 - F(x)$ to denote the complementary cumulative distribution function.

While a user is on the platform, they match at a time-independent exponential rate q, so that over any Δ length time interval starting at any time t, the probability of a user matching is $\Pr(\text{Match} \in [t, t+\Delta)) = 1 - e^{-q\Delta}$. If a user matches we assume they immediately leave the platform.

On the platform's side, we assume the platform commits to some fixed pricing strategy parameterized by a period length, L, and a fixed subscription price, p. We assume users are rational and will pay the subscription price if their expected utility from payment over the subscription period is non-negative. The platform prepares potential matches for the user at marginal operating cost c, paid continuously throughout the user's time on the platform. Finally, we assume the length of time the user can spend on the platform is upper bounded by T, which can be thought of as the size of the pool of potential matches. If a user reaches time T on the platform, they leave, and we say in this case the platform has been *exhausted*.

3.2.1 Model Motivation and Discussion

In this subsection, we delve into the motivation for our model and the key features of online dating it captures. We emphasize that dating platforms are extremely complex and multifaceted ecosystems, for which it would be impossible to succinctly capture all relevant aspects. Instead, our modeling philosophy is to focus on a few salient features that are important and unique to the RM/OM aspects of dating platform management. To organize our discussion, note the model described above has three main components: a user valuation model, a model for matching, and a model for platform strategy.

In our model of user's valuation, we suppose the valuations decay over time at rate δ. This reflects a fundamental aspect of dating platforms where the value of an online subscription is inherently probabilistic but the ultimate outcomes for users are binary, either they match or they do not. Thus, every additional moment a user spends on the platform is one in which they are not obtaining the relationship they are paying to achieve. Moreover, online dating is an arduous process [126] with an extremely high attrition rate. As many four of five users experience emotional burnout or fatigue [104], and survey of 690 online daters who had deleted dating apps [122] found that the number one reason was a need to "take a break" from dating. The parameter δ captures this frustration/fatigue/burn-out as a decay in the users' willingness-to-pay.

In our model of the matching process, we suppose users match at a time-invariant exponential rate q. To build intuition for our model of matching, consider an analogous discrete matching process i.e., one where the user interacts with candidate match 1, then candidate match 2, and so on (see Fig. 8 for a graphical depiction). If the probability of matching with each candidate is q, then a user's time until matching is geometrically distributed with expectation $1/q$. In this chapter we consider the continuous analog of this process indexed by time t, with homogeneous match rate q, such that, over any period of time Δ on the platform, the probability of a user matching is an exponential random variable with rate q, *i.e.*, $\Pr(\text{Match} \in [t, t+\Delta)) = 1 - e^{-q\Delta}$. This formulation is the natural continuous time model for the operation of swipe apps like Tinder, Hinge, or Bumble, where users can continuously swipe through platform-prepared candidate matches.

Figure 8: Graphical depiction of the matching process.

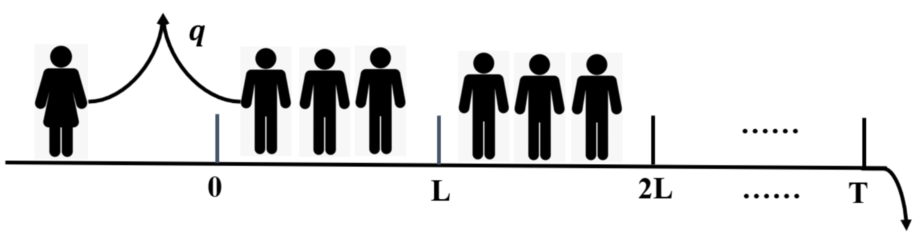

Note. Depicted is the matching process we consider in this chapter. A user interacts with a (continuous) stream of candidate matches. The user forms a successful match with someone in this stream at rate q. The candidate matches are grouped in bunches of size L. In practice, there is a natural correspondence between the length of time L for a period, and the number of candidates a user can see, stemming from hard caps on the number of users who can be seen in a period imposed by the platform.

Finally, our model for the platform strategy supposes the space of policies is parameterized by a subscription price, p, and a subscription length, L. This is a simplification of current app practice where menus of subscriptions are offered, but one we make so as to highlight the differences between these policies. We suppose operating the platform for one additional user has marginal cost c, and note that for large, established platforms with millions of users, c should be vanishingly small. We capture the size of the potential user-base with the parameter T, and note that intuitively when c is small, the platform must be large, and thus T should also be large.

3.2.2 General Profit Derivations

In this subsection, we formally derive the expressions for the general period length pricing strategies. We call a pricing strategy offering a period L of service at a fixed subscription price p, a *period L pricing* (LP).

Suppose the platform commits to a subscription price p with period length L, then each user pays the subscription price at the beginning of every subscription period until they either: 1) match, 2) their expected utility for another period on the platform decays below p, and they decide not to renew their subscription, or 3) the platform is exhausted. Let $\mathcal{R}(p, L)$ be the expected profit for the platform using such a policy where the expectation is taken over the user's valuation, and the matching randomness, i.e.,

$$\mathcal{R}(p, L) := p\,\mathbb{E}\left[\#\text{ Periods on platform}\right] - c\,\mathbb{E}\left[\text{Time on platform}\right], \tag{3}$$

where c is the unit time operating cost. We emphasize the distinction between the number of periods, which is discrete, and the total time, which can include the intra-period time to match. Let $\mathcal{R}(L) := \max_p \mathcal{R}(p, L)$ denote the maximum achievable profit of the platform with payment period of length L.

To compute this profit explicitly, consider a user with fixed valuation v, and note that such a user will pay the subscription price only if their time discounted expected valuation for a match in the subscription period exceeds the price. If the user pays the subscription price, they will stay on the platform until they either match and leave, which is distributed as an exponential random variable with rate q, or until the period ends and they reevaluate paying the period price. Let $\mathrm{TExp}(q, L)$ denote an exponential random variable with rate q truncated at L, which represents the time the user spends on the platform in a period of length L, assuming they paid the period price. A user with valuation $v\delta^{(i-1)L}$ at the beginning of the i^{th} period will pay only if

$$\mathbb{E}\left[\overbrace{v\delta^{(i-1)L+\text{TExp}(q,L)}}^{\text{Valuation for match}}\underbrace{\mathbf{1}_{\text{TExp}(q,L)<L}}_{\text{Chance of matching}}\right]\geq p$$

where $\mathbf{1}_{\text{TExp}(q,L)<L}$ is the indicator function which is 1 when the user matches in the period, and 0 when the user reaches the end of the period without matching. The expectation can be evaluated as,

$$\mathbb{E}[v\delta^{(i-1)L+\text{TExp}(q,L)}\mathbf{1}_{\text{TExp}(q,L)<L}] = v\delta^{(i-1)L}\int_0^L \delta^t q e^{-qt}dt = v\delta^{(i-1)L}\left(\frac{q\left(1-\delta^L e^{-qL}\right)}{q-\log\left(\delta\right)}\right).$$

Therefore, a user with valuation $v\delta^{(i-1)L}$ at the beginning of the i^{th} period will pay the subscription price p only if

$$v\delta^{(i-1)L}\left(\frac{q\left(1-\delta^L e^{-qL}\right)}{q-\log\left(\delta\right)}\right)\geq p. \tag{4}$$

Now we want to express the number of periods the user will purchase, to do so we will need some additional notation. First, let $\omega(i)$ be the minimum valuation necessary to induce purchase in period i, divided by p, i.e.,

$$\omega(i) := \delta^{-(i-1)L}\frac{q-\log\left(\delta\right)}{q\left(1-\delta^L e^{-qL}\right)}. \tag{5}$$

Next, define $\tau(v)$ as the maximum number of periods a user with initial valuation v could possibly pay. When $\tau(v)\leq\left\lfloor\frac{T}{L}\right\rfloor$, $\tau(v)$ satisfies:

$$v\geq\omega(\tau(v))p,\quad v<\omega\left(\tau(v)+1\right)p.$$

Solving the condition for $\tau(v)$ we obtain,

$$\tau(v)=\left(\left\lfloor\left(\log\left(\frac{p}{v}\right)+\log\left(\frac{q-\log\left(\delta\right)}{q\left(1-\delta^L e^{-qL}\right)}\right)\right)/(L\log(\delta))\right\rfloor+1\right)^+,$$

where we use the superscript "+" to denote the maximum with 0.

Now, the number of times a user with valuation v pays the period price p follows a truncated geometric distribution, $\text{TGeo}(1-e^{-qL},\tau(v))$, where $1-e^{-qL}$ is the probability of a successful match in a period, and $\tau(v)$ is the maximum number of periods the user could pay. Thus, the expected number of periods the user with initial valuation v will pay is,

$$\mathbb{E}\left[\text{TGeo}(1-e^{-qL},\tau(v))\right] = \sum_{j=1}^{\tau(v)} j(1-e^{-qL})e^{-(j-1)qL} + \tau(v)e^{-\tau(v)qL} = \frac{1-e^{-\tau(v)qL}}{1-e^{-qL}}.$$

Further, a user with valuation v will stay on the platform at most $\tau(v)$ periods. On the platform, users get matched and leave with rate q. Thus, the time a user with valuation v stays on the platform follows an exponential distribution with rate q truncated at $\tau(v)L$. The expected time the user spends on the platform will be

$$\mathbb{E}[\text{TExp}(q,\tau(v)L)] = \int_0^{\tau(v)L} qte^{-qt}dt + \tau(v)Le^{-\tau(v)L} = \frac{1-e^{-q\tau(v)L}}{q}. \tag{6}$$

Finally, putting it all together the expected profit of the platform can be evaluated as,

$$\begin{aligned}
\mathcal{R}(p,L) &= \int_0^\infty \left(p\mathbb{E}\left[\text{TGeo}(1-e^{-qL},\tau(v))\right] - c\mathbb{E}[\text{TExp}(q,\tau(v)L)]\right) f(v)dv \\
&= \int_0^\infty \left(p\left(\frac{1-e^{-\tau(v)qL}}{1-e^{-qL}}\right) - c\left(\frac{1-e^{-q\tau(v)L}}{q}\right)\right) f(v)dv \\
&= \int_0^\infty \left(p - \frac{c\left(1-e^{-qL}\right)}{q}\right)\left(\frac{1-e^{-\tau(v)qL}}{1-e^{-qL}}\right) f(v)dv \\
&= \left(p - \frac{c\left(1-e^{-qL}\right)}{q}\right)\sum_{i=1}^{\lfloor T/L\rfloor}\left(\frac{1-e^{-iqL}}{1-e^{-qL}}\right)\left(\overline{\text{F}}(\omega(i)p) - \overline{\text{F}}(\omega(i+1)p)\right) \\
&+ \left(p - \frac{c\left(1-e^{-qL}\right)}{q}\right)\left(\frac{1-e^{-\lfloor T/L\rfloor qL}}{1-e^{-qL}}\right)\overline{\text{F}}\left(\omega\left(\lfloor T/L\rfloor+1\right)p\right) \\
&= \left(p - \frac{c\left(1-e^{-qL}\right)}{q}\right)\sum_{i=1}^{\lfloor T/L\rfloor} e^{-(i-1)qL}\overline{\text{F}}(\omega(i)p),
\end{aligned} \tag{7}$$

where the first equality follows from the definition of period L pricing in Eq. (3). The second equality follows from plugging in expressions for $\mathbb{E}\left[\text{TGeo}(1-e^{-qL},\tau(v))\right]$ and $\mathbb{E}[\text{TExp}(q,\tau(v)L)]$. The fourth equality follows from the fact that users with valuation $v \in [\omega(i)p, \omega(i+1)p)$ will all act the same, paying for at most i periods, and $\overline{\text{F}}(\omega(i)p) - \overline{\text{F}}(\omega(i+1)p)$ is the proportion of users who can pay at most i periods, and where the last term $\overline{\text{F}}\left(\omega\left(\lfloor T/L\rfloor+1\right)p\right)$ is the proportion of users whose valuation is larger than $\omega\left(\lfloor T/L\rfloor+1\right)p$. The final equality follows from the fact that $-\left(1-e^{-iqL}\right)+\left(1-e^{-(i+1)qL}\right) = e^{-iqL}\left(1-e^{-qL}\right)$ and simplifying. Note that $p - \frac{c(1-e^{-qL})}{q}$ is the expected profit per period, and $e^{-(i-1)qL}$ the probability that paying users do not match through the first $i-1$ periods.

3.2.3 Profit Derivations for SP and CP

In this subsection, we derive simplified profit expressions for short period pricing, and long period pricing, as limits of the general period L pricing expression in Eq. (7). We will also note some special properties of the profit function for these two cases which will aid in subsequent analysis. Recall, in short period pricing the platform commits to a fixed price p and each user pays the price in a continuous fashion until they either match, or their expected utility from further time on the platform drops below 0, or the platform is exhausted. Let $\mathcal{R}_{SP}(p)$ be the expected profit the matchmaker earns using subscription price p, then,

$$\mathcal{R}_{SP}(p) := (p - c)\,\mathbb{E}\left[\text{Time on platform}\right].$$

We will use $\mathcal{R}_{SP} := \max_p \mathcal{R}_{SP}(p)$ to denote the maximum achievable profit under short period pricing.

In long period pricing, we assume the user agrees to a fixed, one-time payment p, the matchmaker commits to displaying potential matches until the user matches and leaves, or exhausts the platform. Let $R_{CP}(p)$ be the expected profit the matchmaker earns using the long period price p, then,

$$\mathcal{R}_{CP}(p) := (p - c \times \mathbb{E}\left[\text{Time on platform}\right])\,\Pr(\text{User pays price } p).$$

We will use $\mathcal{R}_{CP} := \max_p \mathcal{R}_{CP}(p)$ to denote the maximum achievable profit under long period pricing. The following lemma gives simplified expressions for the profits of SP and CP, respectively.

Lemma 3 (Profit Derivations for SP and CP)**.** *For all positive valued distributions F, and parameters $c, q, T > 0$, and $\delta \in (0,1)$, the expected profit for period L pricing when $L = T$ (CP) or $L = 0$ (SP) is,*

$$\mathcal{R}_{CP}(p) = \left(p - \frac{c(1-e^{-qT})}{q}\right)\overline{F}\left(\frac{p\,(q-\log(\delta))}{q\,(1-\delta^T e^{-qT})}\right) \tag{8}$$

$$= \int_0^T \left(\frac{p\delta^t\,(q-\log(\delta))}{1-\delta^T e^{-qT}} - c\right) e^{-qt}\overline{F}\left(\frac{p\,(q-\log(\delta))}{q\,(1-\delta^T e^{-qT})}\right) dt \tag{9}$$

$$\mathcal{R}_{SP}(p) = \left(\frac{p-c}{q}\right) \int_{\frac{p}{q}}^{\infty} \min\left\{1 - \left(\frac{p}{vq}\right)^{\frac{-q}{\log(\delta)}}, 1 - e^{-qT}\right\} f(v)dv \tag{10}$$

$$= \int_0^T (p-c)e^{-qt}\overline{F}(pq^{-1}\delta^{-t})dt \tag{11}$$

Note, in Lemma 3 two representations are given for both CP and SP. We emphasize that these representations are equivalent, and represent evaluating CP/SP as either an integral over time, or as integral over valuations. Further note, when compared to Eq. (7), these expressions are only in terms of model parameters; there is no complicating sum or function $\omega(\cdot)$. Both expressions for both CP and SP will be useful in our analysis of these policies, hence the four equations above.

Lastly, in this work we will reason about the optimal achievable profit. In describing and analyzing these profits it will be helpful to characterize the profit maximizing prices for a given subscription period L. To ensure these prices are unique, we will often make the ubiquitous assumption that the valuation distribution has a monotone hazard rate (MHR).

Definition 3 (Monotone Hazard Rate (MHR) Distributions)**.** *A random variable $V \sim F$ with density f is MHR if $\frac{\overline{F}(x)}{f(x)}$ is non-increasing.*

MHR distributions are commonly used to model valuations, where they strike the appropriate balance between structure (MHR distributions have sub-exponential tails) and generality. MHR includes many common distributions, including Normal, Uniform, Exponential, and more. In Lemma 4 we show that when valuations are MHR, both the optimal short period and long period prices are unique, and as the cost c increases the prices also increase.

Lemma 4 (Uniqueness of Optimal Prices for SP and CP)**.** *For all positive valued, MHR distributions F, and parameters $c, q, T > 0$, and $\delta \in (0,1)$, the both optimal short period price and long period price are unique, and increasing in c.*

We note that although MHR distributions make the optimal short period and long period prices unique, the same unfortunately can not be said for general period L pricing which can exhibit multiple optimal prices even when the valuations are fixed (see Example 4). Much of the technical difficulty in analyzing general period length L pricing stems from this intractability. We will explore this problem in Section 3.3 when we prove profit guarantees for these pricing policies.

3.3 Profit Analysis

In this section, we study the optimal profits achievable by an online platform in terms of the subscription period length L. To motivate the study of the choice of period length L, let us first emphasize that no fixed choice of L always yields more profit for all parameters and valuation distributions in our setting. In fact, for every choice of L, there is an instance for which that period length is uniquely profit maximizing (see the rightmost panel of Fig. 23 for a construction). With this in mind, in theory, if given parameters c, q, T, δ, and a valuation distribution F, a platform could attempt to directly optimize their choice of L from Eq. (7). However, such optimization is computationally difficult with potentially many locally optimal solutions (see the left panel of Fig. 23). Moreover, the parameters in our model may be difficult to estimate and may vary over time as market dynamics and operational realities change.

Instead of focusing on this explicit optimization problem, we will direct our attention towards proving robust guarantees for some particular choices of L, and on examining natural parameter regimes where the profit maximizing choice of L can be identified. Both approaches will yield results which can directly inform operational decision making with respect to the setting of subscription policies. As mentioned in the introduction, both SP ($L \to 0$) and CP ($L \to T$) will play a starring role as robust solutions, or optimal solutions given some market conditions, respectively. In the next subsection we will show that SP is uniquely robust in the sense that for all market parameters, and any MHR valuation distributions, it achieves a constant factor of the optimal profit against the best-in-class pricing policy. This implies that the potential lost sales from committing to short period pricing versus some unknown optimal policy is never too large.

3.3.1 SP is Always Approximately Optimal

In this subsection, we compare the maximal profit of period L pricing with short period pricing (SP) and demonstrate that SP achieves a constant factor of the profit of the best-in-class policies. This guarantee ensures that the profit of SP is robust to misestimations or fluctuations in the underlying market parameters. Recall, we think of short period pricing as closely resembling a freemium pricing model where the online platform earns revenue from advertising. In advertising supported platforms the user can be thought of as paying with their time/attention in a continuous fashion, and the price rate can be thought of as the intensity/number of ads displayed as they browse the platform. Since almost all major dating platforms in the United States offer a free version of their service (including all Match group apps except `Match.com`), universal profit guarantees for this sort of pricing model are of considerable interest in and of themselves.

Before stating our results, let us first note that SP has several unique properties which make it a prime candidate as a robust pricing policy. For one, by varying the costs c we can construct scenarios for which SP earns positive profit whereas *any* other period length $L > 0$ pricing earns nothing when the costs are prohibitively high (see Example 6 for a construction when valuations are fixed). This immediately precludes any other choice of L from guaranteeing a constant fraction of the optimal profit for all market parameters. Of course, this is just one way of varying the market parameters, in this subsection we will prove that no matter the market conditions, short period pricing always earns at least 26.8% of the profit of any period length L subscription pricing, and for almost all parameters, the guarantee is much stronger.

To prove such a result we will require two lemmas. First, we will show that the optimal profit of any period L pricing can be approximately decomposed into the profit of a short period pricing policy, and the profit of a long period pricing policy where the maximal time on the platform T is reduced to L. Second, we will show that SP approximates CP. Carefully combining these two results will yield our promised robustness guarantee.

Lemma 5 (Profit Decomposition)**.** *For all positive valued distributions F, parameters $c, q, T > 0$, and $\delta \in (0, 1)$, then:*

$$\mathcal{R}(L) \leq \mathcal{R}_{SP} + \mathcal{R}_{CP}^{L},$$

where $\mathcal{R}_{CP}^{L}$ *denotes the optimal profit of long period pricing with* T *set to* L*.*

Lemma 5 decomposes the profit from a period L pricing into the profit from a particular instance of SP and CP. The main idea of the lemma is to note that the profit earned by the end of each subscription period is exactly the same as an appropriately priced instance of short period pricing, and the profit earned in the final subscription period is upper bounded by the profit of a restricted instance of long period pricing. The upside of Lemma 5 is that, while the profit of the optimal best in class period L pricing policy is difficult to write down in closed form and reason about (due to the multiple potential optimal prices), SP and LP are comparatively much easier to compute and analyze. By relating these quantities, we obtain a useful handle on the maximal achievable profit by the best subscription policy (parameterized by L), and demonstrate that reducing our strategy space to just CP and SP maintains approximate optimality.

Armed with Lemma 5, our next task will be to show that SP approximates CP. The following lemma shows that short period pricing always earns at least $1/e$ of the profit of long period pricing for all parameters, and often substantially more.

Lemma 6 ($\mathcal{R}_{SP}$ Approximates $\mathcal{R}_{CP}$)**.** *For all positive valued, MHR distributions* F*, parameters* $c, q, T > 0$*, and* $\delta \in (0,1)$*, then:*

$$\frac{\mathcal{R}_{SP}}{\mathcal{R}_{CP}} \geq \left(1 - \frac{q}{\log(\delta)}\right)^{\frac{\log(\delta)}{q}}.$$

Taking the minimum over $\frac{q}{-\log(\delta)}$ *yields a constant factor approximation,*

$$\frac{\mathcal{R}_{SP}}{\mathcal{R}_{CP}} \geq \frac{1}{e}.$$

Further, this bound is tight.

Lemma 6 gives sharp bounds between SP and CP as a function of the market parameters q and δ. One implication of this result is that for fixed q, and as δ tends to 1, the bound tends to 1, and indeed when $\delta = 1$ these two quantities are equal (see Example 7 for computation). The proof of Lemma 6 is quite involved, proceeding by first reducing the problem to the case of fixed valuations and then constructing a pair of feasible prices of SP that achieve the guarantee over different ranges of the parameter space.

Intuitively, users interact with SP and CP in very different ways. SP gives the most opportunity for users to try the platform since it requires no commitment, however the average time those users spend on the platform is quite short. On the other hand under CP many users are not able to pay the high one time payment and thus do not get access to the platform, but the ones that do pay then stay for the maximal amount time giving them the best chance to get matched. To approximate CP using SP, we consider either a high feasible price for SP that gives a similar opportunity as CP, or a low feasible price that leads to a similar proportion of users getting matched as CP. The best of these two prices achieves the guarantee.

Finally, by taking Lemmas 5 and 6 together, we show that short period pricing can achieve at least $\frac{1}{1+e}$ of the optimal subscription profit of the optimal any period L pricing over any set of market parameters.

Theorem 6 ($\mathcal{R}_{SP}$ Approximates $\mathcal{R}(L)$)**.** *For all positive valued MHR, distributions F, parameters $c, q, T > 0$, $\delta \in (0, 1)$, and for any period length L:*

$$\frac{\mathcal{R}_{SP}}{\mathcal{R}(L)} \geq \frac{1}{1+e},$$

and no other fixed choice of pricing strategy can guarantee any constant factor.

Theorem 6 demonstrates that in all parameter regimes, short period pricing (SP) guarantees a substantial fraction of the profit garnered by period L pricing. We note that our bound is implicitly parametric (via Lemma 6) and often guarantees substantially more than $1/(1 + e)$ of the optimal profit. Thus in practice, we expect SP to earn an even greater fraction of the profit than the worst case guarantee in Theorem 6. Indeed numerically we see that SP is generally quite close to the optimal profit, especially when the costs c and the decay rate δ are fairly high (see Fig. 9 for three examples).

Practically, we think of Theorem 6 as evidence that SP is a robust pricing strategy. This is in-line with its use cases in practice. Often new dating apps are monetized in a short period pricing fashion as they grow their user-base. Unlike long period pricing CP, short period pricing requires no extraordinary market powers or commitments from the users and thus may be the only available pricing mechanism for the emerging online dating platform. These market realities, combined with our guarantee, provide a strong justification for the prevalence of short period pricing for new online dating platforms.

3.3.2 CP is Profit Optimal when Marginal Costs are Low

In Theorem 6, we proved the profit of SP always approximates the profit of an optimally calibrated LP policy. While policies with good guarantees for all markets are desirable, online dating market parameters often are a certain way, and focusing and these more likely parameter ranges may yield stronger guarantees. Specifically, for large and established dating platforms the marginal operation cost c per user should be quite small. Moreover, returning to Example 6 we note that a restriction on c is necessary to give conditions for optimality for any policy besides SP. In Theorem 7 show that such a restriction on the marginal operating cost c is also sufficient for strong profit relations.

Theorem 7 (CP Maximizes Profit when $c = 0$)**.** *For all positive valued, MHR distributions F, and parameters q, T, δ, and any L, if $c = 0$ then,*

$$\mathcal{R}_{CP} \geq \mathcal{R}(L).$$

In Theorem 7 we find that the profit from long period pricing always dominates the profit of *any* other period L pricing when the marginal cost of operating the platform tends to zero. We emphasize that reductions in c improve the profit of all policies regardless of L. To build intuition for this result, recall from our discussion around Lemma 6, the period length L controls a trade-off between opportunity, i.e. how many users get a chance to try the platform, and match proportion, i.e. how many of the users who pay end up actually getting matched. Low L corresponds to higher opportunity, and thus more users on the platform for less time, on average. High L corresponds to less users on the platform for more time. In this context, Theorem 7 states the reduction of cost is more impactful for the smaller cohort of long term users induced by CP, than for any larger cohort of shorter term users induced by another choice of L. So much so that CP becomes profit dominating.

Figure 9: Relations between optimal profit and δ when valuations are uniform.

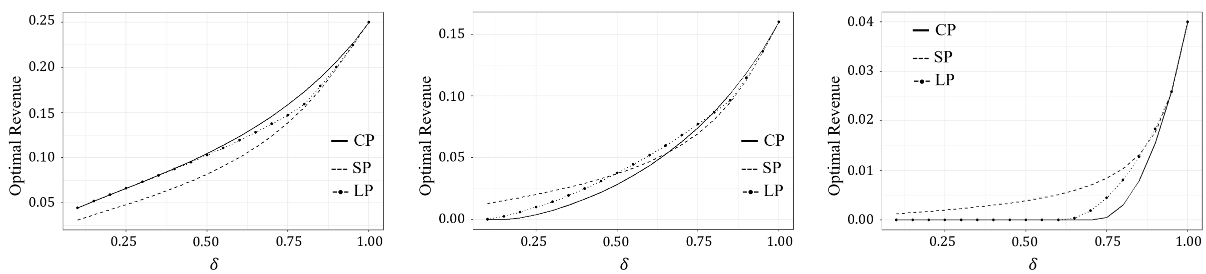

Note. Here we plot the optimal profit under SP, CP and period L pricing with $L = T/15$ for three different levels of marginal cost c. The valuations are drawn from an uniform(0, 1) distribution, the other parameters are $T = 50$, $q = 0.5$, and δ varies. In the left panel, we plot the optimal profits of SP, CP and period L pricing with $L = T/15$ when $c = 0$, and note the profit of CP dominates profits of SP and LP. In the middle panel, we plot the optimal profit of SP, CP and period L pricing with $L = T/15$ when $c = 0.1$, and note profit of SP dominates CP and LP when δ is low, then the dominance switches from SP to LP then CP. In the right panel, we plot the optimal profit of SP, CP and period L pricing with $L = T/15$ when $c = 0.3$, and note the profit of SP dominates CP and LP.

This is in contrast with our previous result which supports SP as the robust choice of pricing strategy. We note that in established online markets one can often expect the marginal cost of operating of the platform to be vanishingly small, and thus although SP is more common for online dating, Theorem 7 suggests that CP may be more profitable for most (established) online dating platforms. Taking Theorems 6 and 7 together, we find the SP and CP have unique profit properties that make them particularly attractive among all period L pricing policies. SP is robust and optimal for emerging platforms with very high marginal costs, and CP is optimal in low cost scenarios which should describe most established online dating platforms.

In Fig. 9, we visualize the optimal profits earned by SP, CP, and an intermediate choice of L, at three different levels of marginal cost c. At each level of cost we vary the user patience parameter δ. We note, as stated in Theorem 7, the profit of CP dominates SP and LP when the marginal cost is vanishingly small, whereas when the marginal cost is high, SP earns more profit than CP and LP, in between the dominance of profit highly depends on δ.

3.4 Match Proportion Analysis

In the previous section we showed that SP and CP have unique properties with respect to profit which make them important to study in comparison with each other. In the remainder of this chapter, we look beyond profit and study the welfare and incentives induced by the choice of SP or CP. In this section we study a specific measure of user welfare, the proportion of the user base that ends up getting matched. In Section 3.5, we extend our study beyond our assumption that the population of candidate matches are all equally likely to match with the user.

3.4.1 Match Proportion and Profit

In this subsection, we will examine the relationship between the profit and the proportion of the market that matches. The proportion of the market that matches is essentially a measure of user welfare, it describes how well a profit maximizing subscription strategy can serve the population that faces it.

First some notation. Let $\mathcal{M}(p, L)$ be the proportion of the market that ultimately gets matched in expectation under period L pricing with price p, and let $\mathcal{M}_{SP}(p)$ and $\mathcal{M}_{CP}(p)$ be the match proportion for SP and CP, respectively. Note, this is the proportion of users who first pay the price, times the probability of them getting matched eventually. As with the profit of different pricing policies, the match proportion can be expressed as a function of the model parameters. For general period L pricing with subscription price p, the probability a user who can pay for at most i periods gets matched eventually is $1 - e^{-iqL}$. The overall match proportion is,

$$\mathcal{M}(p, L) = \sum_{i=1}^{\lfloor T/L \rfloor} \underbrace{\left(1 - e^{-iqL}\right)}_{\text{Prob. of matching}} \underbrace{\left(\overline{\mathrm{F}}(\omega(i)p) - \overline{\mathrm{F}}(\omega(i+1)p)\right)}_{\text{Prob. of paying } p \text{ for } i \text{ periods}}$$

$$+ \underbrace{\left(1 - e^{-\lfloor T/L \rfloor qL}\right) \overline{\mathrm{F}}\left(\omega\left(\lfloor T/L \rfloor + 1\right) p\right)}_{\text{Users who can pay more than } \lfloor T/L \rfloor \text{ periods}}. \tag{12}$$

By taking limits in the same fashion as in Lemma 3, we can derive simplified expressions for the match proportion for SP and CP,

$$\mathcal{M}_{SP}(p) = \int_{\frac{p}{q}}^{\infty} \underbrace{\min\left\{1 - \left(\frac{p}{vq}\right)^{\frac{-q}{\log(\delta)}}, 1 - e^{-qT}\right\}}_{\text{Prob. of matching for valuation } v} f(v)dv, \tag{13}$$

$$\mathcal{M}_{CP}(p) = \underbrace{\left(1 - e^{-qT}\right)}_{\text{Prob. of matching}} \underbrace{\overline{\mathrm{F}}\left(\frac{p\left(q - \log(\delta)\right)}{q\left(1 - \delta^T e^{-qT}\right)}\right)}_{\text{Prob. of paying price } p}, \tag{14}$$

To begin our study of the resulting match proportion from different length pricing policies, we will first note an important relationship between the match proportion, the profit, and the price. Namely, in Lemma 7 we show that the profit of period L pricing is always equal to the match proportion times the difference between *normalized price* and *expected cost*, where the normalized price is $p/\left(1 - e^{-qL}\right)$, i.e., the subscription price for a period over the chance of the user getting matched in the period, and the expected cost is c/q, i.e., the cost times the expected time to match (assuming payment in every period).

Lemma 7 (Profit and Match Proportion)**.** *For all positive valued distributions F, parameters $q, T > 0$, and $\delta \in (0, 1)$, the profit of period L pricing equals the product of the difference between normalized price and expected cost, and match proportion, i.e.,*

$$\mathcal{R}(p, L) = \left(\frac{p}{1 - e^{-qL}} - \frac{c}{q}\right) \mathcal{M}(p, L).$$

Lemma 7 is intuitive, and allows us to leverage some of the insight derived in Section 3.3 to study match proportion. Armed with this lemma, we will investigate whether a larger proportion of the market is matched under SP or CP.

3.4.2 Match Proportion Analysis for SP and CP

Unfortunately as with profit, there is no universal match proportion relationship between SP and CP, or for general period L pricing. Recall when costs are sufficiently high, CP (and any other policy with $L > 0$ by Example 6) is not necessarily economically viable and thus matches none of the user-base, whereas SP may stay in the market and still match at least some users. However, as noted in Section 3.3.2 for online dating one can expect the marginal cost of operating a platform to be relatively small. To rule out these pathological instances and make a fair and relevant comparison between SP and CP, we will assume the cost is vanishingly small.

Theorem 8 (CP Matches More Users than SP)**.** *For all positive valued, MHR distributions F, and parameters q, δ, if $c = 0$ and $T = \infty$, then there exists price p for CP such that:*

$$\mathcal{R}_{CP}(p) \geq \mathcal{R}_{SP} \text{ and } \mathcal{M}_{CP}(p) \geq \left(1 - \frac{\log(\delta)}{q}\right)\mathcal{M}_{SP}.$$

Theorem 8 gives generic conditions for when CP can achieve both a higher profit and match a higher proportion of the market than SP. Thus, in online dating markets with low marginal costs, CP is a win-win: more users are matched, and more profit is made than under SP. At a technical level, the proof is achieved by the serendipitous fact that profit maximizing normalized price for SP also works provably well for CP. We emphasize that such a result does not follow from Theorem 7 and Lemma 7 since the normalized prices can be larger for either SP or CP in certain instances (see the middle panel of Fig. 10).

In Fig. 10, we contextualize Theorem 8 by plotting the relative profits, induced optimal normalized prices, and match proportions as δ varies and $c = 0$. We note that the profit and match proportion of CP dominates SP and LP (for intermediate choice of L) for any δ if the cost is zero. Further, when δ is not too high the induced normalized optimal price of CP is also lower than SP and LP. The numerical results imply CP is indeed a “win-win” pricing policy when the marginal cost is low.

Figure 10: Relations between optimal price, profit, and match proportion when δ changes.

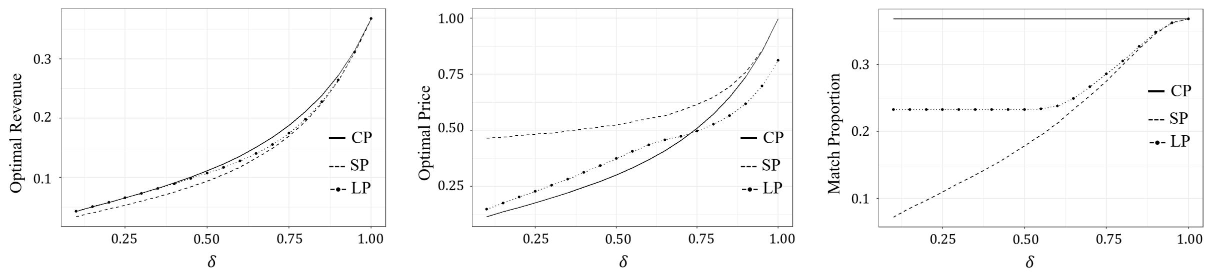

Note. Here we plot the profit, optimal prices, and match proportions under SP, CP and period L pricing with $L = T/15$, when valuations are drawn from an exponential(1) distribution, and where $T = 50$, $c = 0$, $q = 0.3$, and δ varies. In the left panel, we plot the profits of SP, CP and period L pricing with $L = T/15$, and the profit of CP dominates the profits of SP and LP. In the middle panel, we plot the normalized optimal long period, short period and period L prices. In the right panel, we plot the proportion of the market that gets matched under SP, CP and period L pricing, the match proportion of CP always dominates the match proportion of SP and LP.

In Fig. 11 we synthesize Theorems 6 to 8 and plot the relative profits, induced optimal normalized prices, and match proportions as the cost varies. We note that, in these numerics, CP matches significantly more of the market, even when c is not too large suggesting the result in Theorem 8 is relatively conservative. Moreover, the optimal normalized price of CP is often also lower than SP, which means users are paying less to get matched in expectation. Here the results suppose valuations are Exponential, similar results for Uniform valuations and Mixture of Log-normal distributions (which note are not MHR or even unimodal) are shown in Figs. 25 and 26, and similar "win-win" behaviour is exhibited.

Figure 11: Relations between optimal price, profit, and match proportion when valuations are exponential.

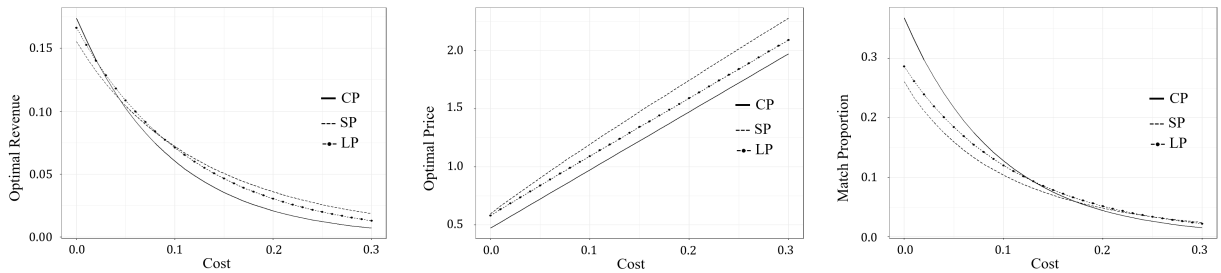

Note. Here we plot the profit, optimal prices, and match proportions under SP, CP and period L pricing with $L = T/7$, when valuations are drawn from an exponential(1) distribution, and where $T = 50$, $\delta = 0.8$, $q = 0.2$, and c varies. In the left panel, we plot the profits of SP, CP and period L pricing with $L = T/7$, and the note relative profit ordering switches from $\mathcal{R}_{CP} > \mathcal{R}_{SP}$ when $c \leq 0.04$, to $\mathcal{R}_{CP} < \mathcal{R}_{SP}$ for $c > 0.04$. In the middle panel, we plot the normalized optimal long period, short period and period L prices. In the right panel, we plot the proportion of the market that gets matched under SP, CP and period L pricing. Note that $\mathcal{M}_{CP}$ dominates $\mathcal{M}_{SP}$ for $c \leq 0.17$.

3.5 Incentive Considerations for Online Dating Platforms

In all previous sections, we have assumed q, the rate at which a user matches on the platform, was fixed and constant. In this section, we will extend our framework to model varying match rates. Many dating platforms, including some of the most popular apps like Hinge, Tinder, Okcupid etc., implement features to segregate high probability and low probability potential matches (e.g. Hinge roses, Tinder top picks, etc.), thus we can infer that users have preferences over potential candidates that impact the match rate, and that platforms are able to learn (to some degree) these match rates and use them to determine the order in which potential matches are displayed.

Formally, let there be k possible match rates $\{q_1, ..., q_k\}$ representing k populations of candidate matches, where $q_1 \leq q_2 \leq ... \leq q_k$, and suppose the online dating platform can order the potential matches based on match rate. Suppose population size of each potential match pool with rate q_i for the user is t_i, $i = 1, ..., k$, where $t_1 + ... + t_k = T$. We assume the user's belief about the match rate on the platform is the average of the potential candidates, $q = \left(\sum_{i=1}^{k} q_i t_i\right)/T$, which is independent of the order in which candidates are displayed. To denote the platform's decision for the order, we use $\overrightarrow{q}$ to denote $\{q_1, ..., q_k\}$, $\overleftarrow{q}$ to denote $\{q_k, ..., q_1\}$, and $\sigma(q)$ to denote any other order of match rates. Let $\mathcal{R}_{SP}(p, \sigma(q))$ and $\mathcal{R}_{CP}(p, \sigma(q))$ to denote the profit of SP and CP, respectively, when potential matches are shown following match rate order $\sigma(q)$ with fixed price p, and we use $\mathcal{R}_{SP}(\sigma(q))$ and $\mathcal{R}_{CP}(\sigma(q))$ to denote the optimal profit of SP and CP under match rate order $\sigma(q)$. For a worked example of how profits change under different orders, see Example 8.

We emphasize that this model for platform controlled match rate order is *not* game theoretic. In our model users cannot learn or respond to a platform's choice of order, and valuations and beliefs about the match rate do not vary as users gain experience. For this reason, we certainly do not claim that the optimal ordering in this model necessarily translates to equilibrium behaviours for the platform. Instead, the purpose of our heterogeneous match rate model is the understand the *obvious* incentives of the online dating platforms with respect to match rate, around which much of the discussion about online dating platforms is centered, and which is often referenced as a reason to innovate in the way dating apps are monetized (cf. Fig. 24). The upshot of our analysis is that if the platform is incentivized to show the best possible matches up front, then the platform and users are in harmony and the space for strategic behaviour collapses.

In our main result of this section, we characterize how a profit maximizing platform orders potential matches under SP and CP, respectively.

Theorem 9 (Incentives for CP and SP)**.** *For all positive valued, MHR distributions F, parameters $c, T > 0$, and $\delta \in (0, 1)$, set of match rates $\{q_1, \ldots, q_k\}$, and for any ordering σ:*

(a) $\mathcal{R}_{CP}(p, \overleftarrow{q}) \geq \mathcal{R}_{CP}(p, \sigma(q))$,

(b) $\mathcal{R}_{SP}(p, \overrightarrow{q}) \geq \mathcal{R}_{SP}(p, \sigma(q))$.

(c) $\arg\max_p \mathcal{R}_{CP}(p, \overleftarrow{q}) \leq \arg\max_p \mathcal{R}_{CP}(p, \sigma(q))$

Theorem 9 shows that a profit maximizing online dating platform has the incentive to show low match rate candidates first under SP, whereas it has the incentive to show high match rate candidates first under CP. This is intuitive, for SP the longer users stay on the platform, the more profit an online dating platform can achieve. In contrast, for CP users pay the long period price p once and thus the strategy to maximize profit is to reduce the operating cost. Therefore online dating platforms using CP want to display high match rates candidates first so that users can get matched and leave the platform as soon as possible.

Interestingly, Theorem 9(c) shows that strategic behavior under CP is not only beneficial for online dating platforms, but also improves users' surplus. The optimal long period price p will be smallest under the mutually beneficial, optimal match rate ordering, and thus under CP the online dating platform can serve even more users by offering a lower price.

More generally, on dating apps, as the user interacts with the platform, the platform learns which potential matches the user would likely prefer. As mentioned in the introduction, ostensibly SP incentivizes the platform to hold likely matches back from the user, whereas CP incentivizes the platform to try and match the user as soon as possible. Taken together, Theorem 9 formalizes this intuition and further shows that not only does CP incentivize the platform to use information about user preferences to help the user match, but also that information induces a profit maximizing platform to lower the long period price. This significantly simplifies the strategic considerations around revealing preference since the user and the platform are, in this regard, perfectly aligned.

3.6 Conclusions

In this chapter we propose a novel and natural model to study the profit obtained by an online dating platform committing to subscription based pricing policies of varying lengths. In our model we gave a number of theoretical guarantees regarding the profit, welfare, and incentive issues induced by the pricing policy which translate directly to managerial insights for future dating platforms. In this final section we will elaborate on some of these managerial insights, and how they relate to current and future dating app operations. We will also touch on some drawbacks of our work, and future directions for research on these critical platforms.

For dating app users and designers, our work sheds light on the consequences of the pricing strategy a profit-maximizing dating apps commits to. As we emphasized in the introduction, a savvy user may recognize that for short subscription period lengths, there is an incentive for any online dating platform to act against their users' best interests so as to earn additional subscription payments. Noting this compatibility mismatch, the strategic user may choose the cheaper option when it would be otherwise utility maximizing to purchase more expensive options. The strategic user may also misrepresent their true preferences, perhaps instead passing themselves off as a casual dater, further complicating the difficult task of finding a match for the user.

This difficult incentive mismatch is referred to as the"strategy puzzle" of online dating by [134]. One way of reading our work is as demonstrating that longer periods offer a way out of this puzzle. Specifically, long period pricing flips this script and aligns the platform and the user. When the costs are paid upfront, the platform and the users have the same goals and the incentive to behave strategically vanishes. We further note that the disincentive to strategize under long period pricing may have additional beneficial effects for the platform that are not accounted for by our model. When all users truthfully report preferences, this in turn may increase the accuracy of the platform's matching algorithms, and thus may generate more matches, which leads to better efficiency, better word of mouth advertising from happy couples, higher return rates, less burn-out among users and so on.

More generally, our work provides theoretical justification for longer periods as a true win-win in the online dating space. When marginal costs are low (as is often the case online), we prove that long period pricing is simultaneously more profitable for the online dating platform, and more effective at matching a significant portion of the user-base. Of course, there are certain barriers not captured in our model that may make longer period pricing strategies like CP more difficult to implement in an online environment. For instance, requiring a large upfront sum from the customer may be difficult for budget constrained users, and requires substantial commitment and trust on the part of the user that after payment the site will be effective for them.

While these implementation issues are potentially burdensome, we believe they are worth the effort of surmounting. We note one emerging avenue for implementing CP is by working in collaboration with government agencies that can track marriage records, and thus enforce contracts where users pay *after* matching/marrying (or never charging them all). In fact, given the strong social value of efficient matchmaking, nationalized dating apps are being tested overseas in countries like Japan and Singapore [5, 2]. For such platforms, CP could reasonably be implemented and may yield superior results.

Finally, we note that online dating in general is remarkably complicated and our model required a number of simplifications to ensure tractability. For instance, our work describes users solely looking for their life-long match, however on sites like Tinder or Grindr, many users may only be looking for short relationships. In this case, the incentives for the site are different, and even under FP the platform may still try to match users in the hopes of soliciting their repeat business. This heterogeneous population of user interests may provide some salve to the analysis presented in this chapter, and would be interesting to study in future work. We also note that our paper studies platforms where a single subscription plan is offered. As seen in Figs. 7 and 24 however, many apps offer a menu of varying length subscriptions, where users can pay for a short, moderate, or long amount of time on the site. In future work, it would be interesting to study how these menus of subscription options effect the incentives and performance of dating apps.

4.0 Fresh Rating Systems: Structure, Incentives, and Fees

4.1 Introduction

Rating systems are the backbone of online marketplaces, facilitating trust between customers and service providers [125]. Customers depend on rating systems to form expectations about service providers, and to avoid undesirable purchasing outcomes. The vast majority of online shoppers consider provider ratings before making purchasing decisions, however ratings alone are often insufficient signals of quality. Instead of relying just on the rating, as many as 88% of online shoppers will take time to further investigate the individual reviews that inform the ratings [129]. This is because not all reviews, and thus not all ratings which are an aggregation of the reviews, are equally informative. When customers investigate reviews they pay special attention to a number of factors that signal how accurate the reviews may be [98].

One key factor customers consider when examining reviews is how long ago the review was posted. In a [129] study they found that 38% of customers will not purchase a product if all reviews for that product are older than 90 days, and 85% of customers report finding reviews older than 90 days to be irrelevant [139]. Intuitively, older reviews are less relevant than newer reviews for purchasing decisions that customers must make today. A review that is sufficiently old so as to be irrelevant to a customer is referred as *stale* [26], and similarly a rating is stale if it is substantially informed by stale reviews. In this work we study how to modify rating systems to mitigate issues due to staling and keep the ratings *fresh*. We make a distinction between two types of stale reviews/ratings which we refer to as sequentially stale and temporally stale. A *sequentially stale* review is one for which there are many more recent reviews. A *temporally stale* review is one which was written a long time ago. As an example, the most recent review of a service provider would be not sequentially stale, but if that review was written six months ago it would be temporally stale. Alternatively if a flurry of reviews were written over a short period (perhaps due to a new product launch or a recent renovation) then the reviews written before this large influx would be sequentially stale even if they written only a short time ago.

In practice, many rating systems treat both stale and fresh (i.e. not stale) reviews equally, regardless of the type of staleness. For instance, consider the popular *simple average* rating system which computes a rating equal to the average of all review scores. The simple average is used by some of the most influential platforms including *Google Reviews* [59] and *Yelp* [142]. However, [26] find that more than half of the reviews on Yelp are stale (defined as older than 90 days). Consequently, Yelp ratings based on simple averages, which do not differentiate between stale and fresh reviews, rely heavily on potentially outdated information. This compromises the credibility of the ratings, a core aspect of Yelp's functionality.

To illustrate this point, consider how staling could result in a rating that misrepresents a service provider's quality. For instance, a high but stale rating could be inflated by a past period of high quality service. Due to the nature of a simple average, if the service provider once performed well but recently experienced a drop in quality, it would require many reviews from this new period of poor service before the rating would reflect the provider's current performance. The converse is also possible, a long flagging service provider that recently changed management or business practices could begin to perform better than it had in the past. However due to the nature of the simple average, the provider would need to acquire at least as many new positive reviews as they had negative reviews when they under-preformed to raise their rating to just the midpoint between the old and new true rating. For examples of such cases in practice see Fig. 30 and Appendix C.3.1.

Another issue with simple average based rating systems is that, for service providers whose rating is already high, there is no incentive to solicit new reviews. Since service providers are often ranked and displayed on platforms based on their ratings, new service providers will expend significant effort to attract positive reviews and build their rating [55]. However once the high rating is established, if there is no penalty associated with the age of the reviews, a service provider may wish to protect their rating by discouraging future reviews [121]. Such behavior is not necessarily malicious, but it inhibits future customers from accessing up-to-date review information reducing trust on the platform.

Stale ratings pose a problem not only for customers but also for the business model of the online platforms that host them. Consider the connection between temporally stale ratings and platform *disintermediation*, which occurs when a customer and service provider connect over a platform, but transact off the platform to avoid platform fees. Disintermediation is a significant concern for online platforms, contributing to the failure of several high profile platforms [67]. Disintermediation occurs when service providers find the value of transacting off-platform higher than the value of transacting on-platform. To connect this to temporally stale ratings, note that many platforms require transaction to occur on platform in order for customers to be able to write reviews. Thus, penalizing reviews for being temporally stale increases the value of each new review to the provider, forcing the service providers to weigh the platform fee against the reputational impact of lost reviews, and potentially encouraging providers to continue transacting on the platform.

In this work we propose and analyze two modifications for rating systems to cope with some of the challenges staling poses. First, we consider a new moving average inspired rating system which systematically down-weights older reviews. We show that within a class of natural rating systems, moving average rating systems optimality protect customers from rating error due to sequentially stale reviews. Second, we study how our moving average system impacts overall platform performance by embedding it into a two-sided ecosystem where the platform facilitates interactions between customers and service providers, including null interactions where service is transacted off-platform and no review is given. In this setting we study the use of penalty reviews which are applied in each null period to shrink high but temporally stale ratings, and incentivize service providers to continuously solicit new reviews. Moreover, since reviews can be locked behind on-platform transactions, we connect this incentive to solicit reviews with disincentives for disintermediating on the platform, and study how the penalty impacts the monetization of the platform.

4.1.1 Our Contributions

The main contributions of our work are as follows:

1. We consider how to design ratings systems to mitigate error due to sequential staling. We introduce and study a general class of *oblivious* rating systems which include common rating schemes such as the simple average. Among all oblivious rating systems, we

prove the asymptotic optimality of the α-moving average rating system in terms of mean squared error over sequences where the quality of service changes (cf. Theorem 10).

2. We compare our proposed α-moving average rating system against popular simple average rating systems and their generalization as L-sliding window rating systems [10]. We theoretically demonstrate that α-moving average rating systems often obtain superior protection against sequential staling (cf. Theorem 11), and numerically demonstrate that error for L-sliding windows is front-loaded whereas the error for α-moving average ratings is smoothly metered out.
3. Finally, we study how α-moving average rating systems can be extended to prevent temporal staling by including penalty reviews. We present a simple model of how market dynamics evolve over time when service providers can choose to transact on or off platform. We characterize the asymptotic behavior of this system (cf. Proposition 1), and give theory to guide platform designers in setting the fee and penalty review term. As a function of the minimum quality allowed on the platform, we give closed-form, prior-free expressions for the fee and penalty that incentivize maximum participation on the platform and guarantee at least $3/4^{th}$ of the optimal revenue for the platform (cf. Theorem 12).

Overall, our work provides tools for platform designers to use to cope with issues due to stale reviews. For the problem of sequential staling, our work makes a case for the α-moving average rating systems as a powerful alternative rating mechanism. In many ways α-moving average systems can be seen as a smooth implementation of L-sliding window policies, inheriting all of the virtues of L-sliding windows while avoiding the sharp cutoff after L periods. For the problem of temporal staling, we advocate for the introduction of penalty reviews that down-weight ratings for inactive service providers and give theory to describe how to set the penalty parameter. One surprising upside of penalty reviews is their connection with disincentivizing disintermediation which we hope will further motivate practitioners to experiment with them in practice.

4.1.2 Literature Review

Given the importance of rating systems for online marketplaces, the scope of academic work that studies them is vast and multidisciplinary. Here we describe some of the most relevant streams of literature, and explain how our work contributes to and/or differs from each.

4.1.2.1 Reputation Systems in Online Marketplaces

Rating systems, sometimes referred to as reputation systems, are critical for facilitating trust in online marketplaces [125, 93]. As such, rating systems have a central place in the theory of two-sided online market places [110, 1]. A number of papers have looked at how ratings impact how customers and service providers participate on the platform. Some recent representative examples include [69] who study how to design rating systems to induce maximum effort from participants, and [118] who highlight the information asymmetry between the platform and its users, and study how platforms should dispense ratings to market participants. Relative to these works and works cited within, our focus is not on what kind of rating information the platform releases, but instead on the mechanics of how a platform should aggregate reviews into ratings so as to prevent staling.

More in line with our work, there is a influential stream of papers that look at how reviews can be faked or altered so as to manipulate ratings. [94] study the massive proliferation of fake reviews in major online platforms like Yelp. [90] looks at how sellers can buy fake reviews, and how platforms should adjust to account for this behavior. Similarly, [97] look at how sellers can bribe reviewers into giving credible but biased reviews for their products. These papers all study ways in which the rating displayed on the platform can fail to reflect the true quality of a service provider. We contribute to this literature by looking at another way ratings can misrepresent service provider quality via staling, and combat it by modifying the structure of the rating system.

4.1.2.2 Design and Analysis Rating Systems

The our works contributes to a stream of literature that considers how exactly ratings should aggregate reviews. There are many design aspects of ratings to consider, for instance [68, 111] consider how ratings should aggregate reviews into discrete buckets corresponding to the number of stars. Our work is closely literature related to the problem of how ratings system should remember old reviews. [83] discuss the impact of deleting old reviews, which is essentially the operational framework of L-sliding window rating systems introduced in [9, 10] and developed in [57]. Our α-moving average rating systems take a similar perspective on ratings, but instead of dropping all but the L most recent reviews, we smoothly degrade the impact of the review on the rating as a function of the review's age.

We emphasize that we are not the first to consider moving average inspired rating systems. [54] introduce a similar rating system, we they call the exponential smoothing mechanism, and analyze how it can induce trustworthy behavior from the seller. [32] also describe a model with ratings aggregated according to a moving average, however their focus is to study the confounding effect of price and rating. Relative to these works, we are one of the first to give serious attention to this rating system as a way of mitigating issues related to staling.

More generally we note that our analysis of moving average inspired rating systems as a way of detecting and responding to change in the underlying service quality distribution is similar to some work on the problem of *change point detection* [109, 92, 91]. The focus in these works is on identifying the moment when a change occurs, whereas in our work our focus is on how much error is incurred by such change as the ratings adjust. As such, much of the modern work on change point detection has used dynamic statistically tests to flag changes in distribution [7], whereas in our work we preclude such approaches by focusing on oblivious rating systems which are comparatively simpler to understand and implement.

4.1.2.3 Disintermediation on Online Platforms

In the latter half of this work we consider a novel connection between temporal staling and market disintermediation. Market disintermediation has been a subject of intense interest, first studied in the context of supply chains [108, 3], but recently as major source of market failure for online platforms [18]. Market disintermediation has been identified as a cause of several high-profile online marketplace collapses [79, 88, 140]. In response, a stream of papers in business strategy [63, 136] and operations management [66, 116] have looked at models to capture the causes of such disintermediation and proposed a variety of structural and price based remedies. Our work is, to the best of our knowledge, the first to explicitly study how down-weighting old ratings can incentivize service providers to transact on the platform, discouraging disintermediation.

4.1.3 Chapter Outline

The remainder of this chapter is organized as follows. In Section 4.2, we consider moving average inspired rating systems and demonstrate that they offer excellent protection against sequential staling. In Section 4.3, we situate the recommendation system in a general model for transactions on two-sided platforms and characterize the equilibrium market outcomes in the case where the platform can give penalty reviews to down-weight temporally stale ratings. In Section 4.4, we study prior-free ways to set the penalty review term and the platform's fee to (approximately) maximize platform revenue. Finally, in Section 4.5 we conclude and highlight avenues for future work.

All examples and proofs referenced in the main body can be found in Appendices C.1 and C.2 in the Appendix. Code for all numerics in the chapter is publicly available at `https://github.com/tcui-pitt/Ratings`

4.2 Sequential Staling and Moving Average Rating Systems

In this section, we initiate our formal study of *rating systems* and consider the question of choosing a rating system to mitigate issues due to sequential staling.

4.2.1 Model and Preliminary Definitions

We consider reviews as being generated one at a time, where time is indexed by t. Each review represents a customer's experience with the service provider and is modeled as a positive-valued noisy signal, s_t, which is drawn independently from some underlying distribution of service provider quality in that time period, F_t. We assume F_t always admits a density f_t, and has mean quality μ_t and standard deviation σ_t. The rating system aggregates these reviews into a numeric score. Formally,

Definition 4 (Rating System). *A **rating system** is a sequence of functions $\{R_t\}_{t=1}^{\infty}$ that map from the sequence of reviews $\{s_t\}_{t=1}^{\infty}$ into a numeric score between 0 and 1, i.e., where $R_t : \{s_i\}_{i=1}^{t} \to [0,1]$.*

This is a generic definition of rating systems and captures all common rating systems. We highlight two systems of particular relevance that were mentioned in the introduction: the *simple average* rating system is such that the rating at time t is the average of reviews, i.e., $R_t^{SA}(\{s_i\}) := \frac{\sum_{i=1}^{t} s_i}{t}$. The *L-sliding window* rating system parameterized by L, the width of the window, is such that the rating at time t is the average of the L most recent reviews, i.e., $R_t^{SW}(\{s_i\}) := \frac{\sum_{i=(t-L)^+}^{t} s_i}{\min\{L,t\}}$ [1] where the notation $(x)^+$ denotes $\max\{x, 0\}$. Note that without loss of generality we normalize ratings to fall between 0 and 1 as opposed to common implementations such as the star ratings out of five.

Throughout this section we will give special attention to a lesser known system that we term the *α-moving average* rating system [2]. The α-moving average rating system is defined recursively as,

$$R_t^{MA}(\{s_i\}_{i=1}^{t}) = (1-\alpha)R_{t-1}(\{s_i\}_{i=1}^{t-1}) + \alpha s_t \text{ for } t > 1, \quad R_1 = s_1. \tag{15}$$

In the α-moving average rating system, the rating at time t is a weighted average of the previous rating (i.e. the rating at time $t-1$) and the most recent review. As with the L-sliding window rating system, the α-moving average rating system depends on a parameter,

[1]See [10] for more on this rating system.

[2][54] refer to this system as the *exponential smoothing* rating mechanism. We prefer to refer to it as a moving average in parallel with the simple average, and following the naming convention in time series analysis which prominently feature aggregation mechanisms of a similar form.

$\alpha \in (0, 1)$, which determines the weight the system places on the most recent review. When the rating system being discussed is clear from context we will omit the superscripts.

Note that our definition of rating systems as a sequence of functions is extremely broad, and does not enforce many features commonly associated with rating systems. For instance, in our definition there is no notion of correctness. To narrow down our search for a rating system that can mitigate issues due to sequential staling, we will restrict our attention to the natural class of rating systems that can be described by a fixed sequence of weights. Such rating systems have a predetermined way of aggregating reviews over time and are always correct in expectation when the underlying quality distribution F_t is not changing over time (i.e. is stationary). Since the weights are predetermined, these systems are not adapted to the exact values of the reviews that appear. As rating systems in this class are not adaptive, we call these types of rating systems *oblivious*. Formally,

Definition 5 (Oblivious Rating Systems)**.** *We call a rating system* ***oblivious*** *if there exists a predetermined sequence of weights* $\{w_i\}_{i=1}^{\infty}$ *such that at any time* t*, the rating is equal to*

$$R_t(\{s_i\}_{i=1}^t) := \frac{\sum_{i=1}^t w_i s_i}{\sum_{i=1}^t w_i}. \tag{16}$$

Oblivious rating systems are a natural class of rating systems to study. As they are non-adaptive, oblivious rating systems of this form are simple and easy to implement, making them attractive to practitioners, while still being quite expressive. The simple average rating system is oblivious with all weights being identical (i.e. $w_i = w_j$ for all i, j). The α-moving average rating system is also oblivious with the weights forming a geometric sequence ($w_1 = 1, w_i = \alpha/(1-\alpha)^{i-1}$ for $i \geq 2$). However not all common rating systems are oblivious, namely the L-sliding window rating system is not oblivious since it dynamically changes the weight it gives to a review to zero when that review leaves the width L window.

In addition to being simple to describe, oblivious rating systems have many desirable properties. They are always correct in expectation when signals are generated IID, i.e. when $F_i = F_j$ for all i, j the expected rating at every time t equals the mean review by linearity of expectation. This property, commonly defined as *consistency*, is important since the purpose of a rating system is that it gives a reliable estimate of a provider's true quality of service. We are now ready to begin our search for a rating system that is robust to issues caused by sequential staling. Namely, among the class of oblivious rating systems we wish to determine which one offers the best customer protection against sequential staling.

4.2.2 Fresh Oblivious Rating Systems

Recall from our discussion in Section 4.1 that the key way sequential staling introduces error into the ratings is by preventing the rating from detecting and adapting to changes in a service providers quality. When abrupt change in the underlying service quality occurs, previous reviews about a provider become stale, and ratings based on those reviews are no longer an accurate predictor of service quality - often to the detriment of new customers (for example recall Fig. 30). We are interested in finding rating systems that are malleable enough that when change occurs, the rating can quickly adjust and become, again, informative. Intuitively then, a rating system that is not vulnerable to sequential staling is one that gives a rating which closely tracks with the most recent reviews. Of course the rating system can not only give weight to the most recent review since such a rating will then be very noisy and incur high expected error (measured as the squared difference between the rating and the current expected service quality) when a change does not occur.

Mathematically, we can translate this notion as requiring that the rating system should balance two types of rating error: the error which occurs when the underlying distribution of service quality changes or *switches*, and the error which occurs when it does not. To study this we will consider the problem of minimizing the maximum of these two sources of rating error incurred over sequences of service quality distributions with the mechanics that, at some a priori *unknown* time t, the distribution of service quality can switch. Formally, we will fix distributions F_1 and F_2 with the mechanics that in period i for $i < t$, the review $s_i \sim F_1$ and for $i \geq t$, the review is now $s_i \sim F_2$. These sequences capture issues due to sequential staling described above, with the switch from F_1 to F_2 representing a fundamental change in provider service quality. We emphasize that, for the oblivious rating system, since this switching time t is unknown, the weights $\{w_i\}$ that define the rating must be chosen in such a way that the sequence is adaptable no matter when the switch occurs.

A natural question then is what oblivious rating system can best minimize the maximum of these two sources of error? In the main theorem of this section, we show that if a platform chooses the weights of their rating system to minimize the maximum rating error incurred either by a switch in the underlying provider quality, or by the variance in the ratings when no switch occurs, this rating system will always be asymptotically equivalent to an α-moving average rating system. As balancing these two forms of error at each time step effectively inoculates the rating system against sequential staling, we conclude that the moving systems are ideal for tackling issues due to staling.

Theorem 10 (Asymptotic Optimality of Moving Average Ratings)**.** *Consider any sequence of distributions defined by F_1 with mean μ_1 and standard deviation σ_1, and F_2 with mean $\mu_2 \neq \mu_1$ and standard deviation σ_2, where at any time a switch from F_1 to F_2 could occur. If the weights w_t of an oblivious rating system satisfy:*

$$\min_{w_t}\left\{\max\left\{\underbrace{\mathbb{E}\left[\left(\frac{\sum_{i=1}^t w_i s_i}{\sum_{i=1}^t w_i}-\mu_1\right)^2 \middle| \{s_i\}_{i=1}^t \sim F_1\right]}_{\text{No switch error}}, \underbrace{\mathbb{E}\left[\left(\frac{\sum_{i=1}^t w_i s_i}{\sum_{i=1}^t w_i}-\mu_2\right)^2 \middle| \{s_i\}_{i=1}^{t-1} \sim F_1, s_t \sim F_2\right]}_{\text{Switch error}}\right\}\right\},$$

for every t, then there exists some $\alpha \in [0,1]$ such that

$$\lim_{t\to\infty}\frac{w_t}{\sum_{i=1}^t w_i}=\alpha,$$

and thus $\{w_i\}_{i=1}^{\infty}$ is asymptotically equivalent to an α-moving average rating system.

Theorem 10 identifies the α-moving average system as the asymptotically optimal oblivious rating system for protection against sequential staling. Let us take a moment to unpack the result stated here: at each period t a switch in the distribution of service quality from F_1 to F_2 could occur. If it does not occur, then the "no-switch" error is the MSE for the rating system versus the true mean quality at t which is μ_1. If a switch does occur, then the "switch" error is the MSE of the rating system versus the new true mean quality at t which is μ_2. In Theorem 10, we suppose the sequence of weights are chosen so that at time t, w_t minimizes the maximum of these two sources of error which is a natural way to choose the w_t when the exact time of the switch is unknown. Note then that the sequence of weights $\{w_i\}_{i=1}^t$ is defined iteratively: $w_1 = 1$ without loss of generality, then w_2 minimizes the maximum error modulo the fixed value of w_1, then w_3 minimizes the maximum error modulo the values of w_1 and w_2, and so on. The guarantee in Theorem 10 then is that for a sequence of weights built in this way, in limit the value of w_t is proportional to the sum of all weights, $\sum_{i=1}^{t} w_i$, which is equivalent to becoming an an α-moving average system with $\alpha = \lim_{t\to\infty} \frac{w_t}{\sum_{i=1}^{t} w_i}$ since for large t,

$$\begin{aligned}
R_t(\{s_i\}_{i=1}^t) &= \frac{\sum_{i=1}^{t} w_i s_i}{\sum_{i=1}^{t} w_i} \\
&= \frac{\sum_{i=1}^{t-1} w_i s_i}{\sum_{i=1}^{t} w_i} + \left(\frac{w_t}{\sum_{i=1}^{t} w_i}\right) s_t \\
&= \frac{\sum_{i=1}^{t-1} w_i s_i / \sum_{i=1}^{t-1} w_i}{\sum_{i=1}^{t} w_i / \sum_{i=1}^{t-1} w_i} + \left(\frac{w_t}{\sum_{i=1}^{t} w_i}\right) s_t \\
&= R_{t-1}(\{s_i\}_{i=1}^{t-1}) \left(\sum_{i=1}^{t-1} w_i / \sum_{i=1}^{t} w_i\right) + \left(\frac{w_t}{\sum_{i=1}^{t} w_i}\right) s_t \\
&= \left(1 - \frac{w_t}{\sum_{i=1}^{t} w_i}\right) R_{t-1}(\{s_i\}_{i=1}^{t-1}) + \left(\frac{w_t}{\sum_{i=1}^{t} w_i}\right) s_t \\
&\approx (1-\alpha) R_{t-1}(\{s_i\}_{i=1}^{t-1}) + \alpha s_t,
\end{aligned}$$

where the third equality follows from dividing $\sum_{i=1}^{t-1} w_i$ for both the numerator and denominator in the first term, and the final approximation follows from the guarantee in Theorem 10.

The proof of Theorem 10 is relatively involved, but the intuition for the result is simple. There are two ways the rating system can make error, either from the noise in the rating system when no switch occurs, or from miscalibration due to weight put on older reviews when a switch does occur. Thus, to minimize the two types of error the system should put a sufficient fraction of the weight on the most recent review, intuitively keeping the two forms of error in balance. Since the switching time at which the change in underlying service quality is unknown to the rating system, and as an oblivious rating system's weights must be specified at the beginning of the period, the rating system must act as if every at time t a change can occur. As such, in limit it always puts a constant weight on the most recent observation and thus automatically is a moving average rating system. There are two additional technical difficulties that must be overcome to complete the proof. The first difficulty is that the error before the switch is not constant but instead slowly decaying as more information is fed into the system. In practice this error is minimal (as we will demonstrate in Section 4.2.4) and here is handled by the asymptotic nature of the result. The second difficulty is that, even asymptotically, when the noise in the switching distribution is large (i.e. $\sigma_1 < \sigma_2$) these two forms of rating error should not always be made equal. Instead it is sometimes optimal to purely attempt to minimize the switching error, since the no-switch is always going to be relatively small. The problem is further complicated by "pooling" effects where, when μ_1 is close to μ_2 but σ_2 is much larger than σ_1, then keeping weight on previous stale reviews is useful to the rating system since the stale reviews are approximately correct but less noisy. These effects together influence the optimal choice of the parameter α so that it depends on the moments of F_1 and F_2 in quite intricate ways which require an involved analysis.

In terms of the practical implementation of moving-average rating systems, it should be noted that the guarantee in Theorem 10 requires the system to have upfront knowledge of the moments of both distributions, μ_1, μ_2 and σ_1, σ_2, in order to appropriately set the weights $\{w_i\}$. In practice, the first set of moments can be implemented as the average review and its variance among all providers on the platform and thought of as a sort of baseline. The second distribution, F_2, then acts as design lever for the rating platform. The assumed moments of F_2 represents a level of desired "safety" with respect to the type of outcomes the platform thinks is possible and wants to protect their customers from. We emphasize that the exact forms of F_1 and F_2 are irrelevant for the theorem, it depends only on the first two moments of those distributions.

Finally, based on the set levels for F_1 and F_2, Theorem 10 shows the optimal weights converge to an α-moving average rating system. One natural question is that, if these optimal weights are computable, why not use them directly instead of reasoning about their limiting properties? To this end, we emphasize that the weights are quite complicated whereas the moving average is simpler to understand both for the platform and its users. We also note that the convergence of the optimal weights to a moving average system is quite fast, see the right panel of Fig. 28 for an example. Finally, the moving average is much easier to reason about than the optimal weights, involving just a single parameter, α. While Theorem 10 implies the existence of an optimal α for a moving average system, it does not actually solve for the best choice of the parameter. Indeed, the optimal specification for α can often be written down in closed-form, which we describe next.

Corollary 1 (Choice of α)**.** *Consider any sequence of distributions defined by F_1 with mean μ_1 and standard deviation σ_1, and F_2 with mean μ_2 and standard deviation σ_2, where at any time a switch from F_1 to F_2 could occur. Suppose $|\mu_1 - \mu_2| \geq \frac{\sigma_1}{2}$ and $\sigma_1 > \sigma_2$, and let $\alpha^* = |\mu_1 - \mu_2| / \left(|\mu_1 - \mu_2| + \sqrt{\sigma_1^2 - \sigma_2^2}\right)$. Then the choice of α for an α-moving average system that asymptotically minimizes the objective,*

$$\min_{\alpha} \lim_{t\to\infty} \left\{ \max \left\{ \mathbb{E}\left[\left(R_t^{\alpha\text{-}MA} - \mu_1\right)^2 \middle| \{s_i\}_{i=1}^{t} \sim F_1 \right], \mathbb{E}\left[\left((R_t^{\alpha\text{-}MA} - \mu_2\right)^2 \middle| \{s_i\}_{i=1}^{t-1} \sim F_1, s_t \sim F_2 \right] \right\} \right\},$$

is:

(a) If $2\left(\frac{\sigma_1^2}{(2-\alpha^*)^2} - (\mu_1 - \mu_2)^2(1-\alpha^*) + (\sigma_2^2 - \sigma_1^2)\alpha^*\right) \leq 0$, the choice of α that minimizes the maximum MSE in limit is:

$$\alpha = \alpha^*.$$

(b) If $2\left(\frac{\sigma_1^2}{(2-\alpha^*)^2} - (\mu_1 - \mu_2)^2(1-\alpha^*) + (\sigma_2^2 - \sigma_1^2)\alpha^*\right) > 0$, the choice of α that minimizes the maximum MSE in limit is the unique solution to,

$$2\left(\frac{\sigma_1^2}{(2-\alpha)^2} - (1-\alpha)(\mu_1 - \mu_2)^2 + \alpha(\sigma_2^2 - \sigma_1^2)\right) = 0,$$

such that $\alpha \in [0, 1]$.

Corollary 1 gives the optimal choice of α as a function of the starting distribution and the switching distribution for minimizing the maximum error objective in Theorem 10. Corollary 1 requires a few conditions in order to express the value of α in (implicit) closed form. We emphasize that an optimal choice of α always exists by Theorem 10, however when μ_1 and μ_2 are not sufficiently different i.e., $|\mu_1 - \mu_2| < \frac{\sigma_1}{2}$, then the optimal choice of α converges to 0. Moreover, when $\sigma_1 \leq \sigma_2$, the switch error $\mathbb{E}\left[\left((R_t^{\alpha\text{-MA}} - \mu_2\right)^2 \middle| \{s_i\}_{i=1}^{t-1} \sim F_1, s_t \sim F_2 \right]$ always dominates, and the choice of α is can be obtained by solving the first order condition to minimize the switch error.

In more practical terms, when the concern about sequential staling is high, the value of α should increase which in turn increases the weight placed on the most recent review. In this way, α is tunable knob for the platform that determines the flexibility/malleability of the system. However, increasing α also increases the variance in the ratings which can lead to error when the service distribution is stable. It will often be useful to speak directly about the variance of the ratings in a system, which we capture in the following definition.

Definition 6 (c-Consistency)**.** *A rating system is **c-consistent** if for any distribution F such that signals $s_t \sim F$ are drawn IID with finite mean μ and variance σ^2, the asymptotic rating is such that $\lim_{t\to\infty} Var[R_t] \leq c\sigma^2$.*

We will refer to a rating system as *exactly c-consistent* if the consistency condition holds with equality i.e. $\lim_{t\to\infty} \mathrm{Var}[R_t] = c\sigma^2$. Note by definition of c-consistency, we always assume $c \leq 1$. In the next subsection, we will compare the α-moving average rating and the L-sliding window rating systems when both of them are exactly c-consistent.

4.2.3 Moving Average vs. Sliding Window Rating Systems

In the previous subsection we identified the α-moving-average rating system as the myopically optimal oblivious rating systems for protecting against change in the underlying quality distribution. While oblivious rating systems capture many familiar implementations, L-sliding window rating systems are not oblivious since they dynamically change the weight of ratings when they leave the window. In this subsection, we will directly compare α-moving average and L-sliding window rating systems. At high level, both moving average and sliding window systems extend simple average systems by systematically down-weighting older reviews. Thus in a sense both system address the problem of sequential staling. The difference is that while the L-sliding window system simply drops sufficiently old reviews, the α-moving average rating system does this down-weighting in a smooth way. In this way, the α-moving average system can be thought of as a smoothed version of an L-sliding window rating system. In the main result of this subsection we formalize this correspondence.

First, to give a fair comparison the α-moving average and L-sliding window as systems that can mitigate sequential staling by balancing error when a switch does/does not occur, we need to set the parameters of the two systems (α and L) in an equivalent way. Since both systems are always correct in the expectation in the case when no switch occurs, the most natural way to put them in correspondence is to look at their second moment and enforce that both systems have the same asymptotic rating variance. Thus, for the remainder of this subsection we will compare the two systems setting α and L so that both systems are exactly c-consistent. In Lemma 8 we write down correspondence between c, α, and L in closed-form.

Lemma 8 (Consistency of Rating Systems)**.** *Given some service quality distribution F with mean μ and variance σ, such that review signals in each period are drawn IID from F.*

(a) To be c-consistent, an α-moving average rating system must have $\alpha \leq \frac{2c}{1+c}$.

(b) To be c-consistent, an L-sliding window rating system must have $L \geq \lceil \frac{1}{c} \rceil$.

Note, unlike α in the α-moving average rating system, the L for the L-sliding window rating must be an integer since it represents the length of the window. For the sake of comparison in this subsection we will always consider c such that $1/c$ is an integer. Now we are ready to compare how c-consistent α-moving average systems and L-sliding window systems process reviews. Specifically, for the objective of balancing switching versus no-switching error as in Theorem 10, we consider which implementation is superior, and when.

Theorem 11 (Sliding Window vs. Moving Average)**.** *Consider any sequence of distributions defined by F_1 with mean μ_1 and standard deviation σ_1, and F_2 with mean μ_2 and standard deviation σ_2 where $\mu_1 \neq \mu_2$, and some time t^* such $F_t = F_1$ if $t \leq t^*$ and $F_t = F_2$ if $t > t^*$. If α and L are chosen so that, for some $c > 0$, both systems are exactly c-consistent, then:*

(a) The MSE of the L-sliding window rating system immediately after service quality switches is

$$\lim_{t^* \to \infty} \mathbb{E}\left[(R^{SW}_{t^*+1} - \mu_2)^2 \middle| \{s_i\}_{i=1}^{t^*} \sim F_1, s_{t^*+1} \sim F_2 \right] = c\sigma_1^2 + (1-c)(\mu_1 - \mu_2)^2 + c^2(\sigma_2^2 - \sigma_1^2).$$

(b) The MSE of the α-moving average rating system immediately after service quality switches is

$$\lim_{t^* \to \infty} \mathbb{E}\left[(R^{MA}_{t^*+1} - \mu_2)^2 \middle| \{s_i\}_{i=1}^{t^*} \sim F_1, s_{t^*+1} \sim F_2 \right] = c\sigma_1^2 + \left(\frac{1-c}{1+c}\right)^2 (\mu_1 - \mu_2)^2 + \left(\frac{2c}{1+c}\right)^2 (\sigma_2^2 - \sigma_1^2).$$

(c) When $\sigma_1^2 \geq \sigma_2^2$, the maximum MSE of the α-moving average rating system over the time horizon is always less than the maximum MSE of the L-sliding window rating system

Theorem 11 compares the mean squared error of equivalently parameterized moving average and sliding window rating systems. Note that when the systems are both exactly c-consistent, the expected no-switching error for both systems is identical - a fact that should be unsurprising given that the no-switching is directly related to the variance of the rating system. What Theorem 11 shows is that the switching error for the moving average rating system is often less than for an equivalent L-sliding window system. Thus, when the switching error dominates the no-switching error, moving average systems offer superior protection against sequential staling. This dominance requires a condition, $\sigma_1^2 \geq \sigma^2$, which means that the variance in the distribution of service quality is lower after the switch.

To prove the result we decompose the after switching MSE into three sources. First is the error from the mismatch between the parts of the rating informed previous reviews drawn from F_1, the second stems from service quality difference between F_1 and F_2 when the service quality switches, and the third is the error from the most current review which depends on the variance of new reviews, σ^2. For the L-sliding window and α-moving average rating system, the error of this first kind from previous reviews are the same since both of them are c-consistent. However, from Lemma 8, we note that the α-moving average rating system puts more weight on the most recent review. Thus this second form of error from the service quality difference for the α-moving average rating system is always less. Finally, for the third source of error we use our condition that the variance of the most recent review is smaller, i.e., $\sigma_2 \leq \sigma_1$, to conclude that the overall MSE of the α-moving average rating system will always be less than the L-sliding window.

Finally, to prove both Theorems 10 and 11 we our analysis required looking at error in limit. In Section 4.2.4 we numerically explore the differences between rating systems to observe how our theoretical results translate to finite settings.

4.2.4 Rating System Numerics

To conclude this section we will numerically explore the differences between simple average, α-moving average, and L-sliding window rating systems, in the setting where reviews that are drawn from a (possibly time-varying) exponential service quality distribution with rate λ truncated to $[0, 1]$. We denote the truncated exponential by $\text{TExp}(\lambda) := \min\{\text{Exp}(\lambda), 1\}$. As the truncated exponential is a single parameter family, we can use it to model changes in the underlying service quality distribution simply by varying the distribution parameter λ. In these numerics, we will consider what the rating error looks like when there is a sudden change in the underlying service quality distribution, and how long it takes the system to return to equilibrium.

First, in Fig. 12 we consider the mean squared error incurred by rating systems when the underlying service changes between $s_t \sim \text{TExp}(2)$ and $s_t \sim \text{TExp}(4)$. This experiment represents a finite version of the setting considered in Theorems 10 and 11 and models the situation where the rating becomes stale due to a sudden change in the underlying quality of the service provider. By studying the rating error we can observe how long it takes for different rating systems to "learn" the new true rating.

Figure 12: MSE as for Various Rating Systems as Service Quality Changes.

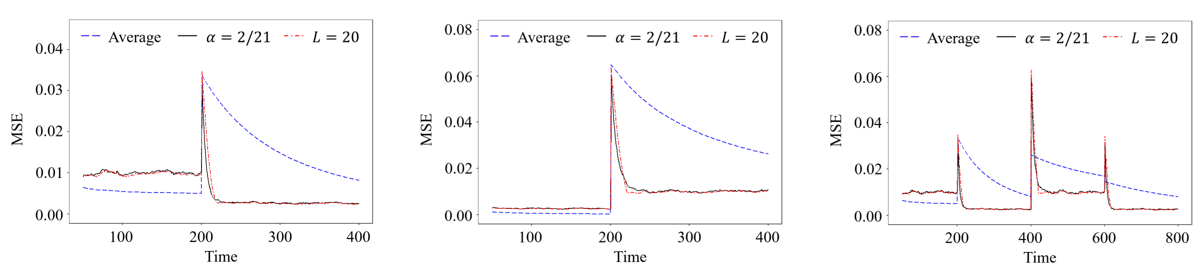

Note. We plot the α-moving average, L-sliding window, and simple average rating system's MSE over time when the reviews for the service provider are drawn from a truncated exponential distribution, and where the service quality distribution changes over time. In the left panel, we plot the simulated MSE over 1000 trials when the underlying service quality distribution changes from $\lambda = 2$ to $\lambda = 4$ at time $t = 200$. In the middle panel, we plot the simulated MSE over 1000 trials when the underlying service quality distribution changes from $\lambda = 4$ to $\lambda = 2$ at time $t = 200$. In the right panel, we plot the simulated MSE over 1000 trials when the underlying service quality distribution changes from $\lambda = 2$ to $\lambda = 4$ and back, with a period length of 200.

From Fig. 12 we can see that both the α-moving average system and L-sliding window system can quickly adapt to changes in the service quality. We emphasize that although Theorem 10 requires asymptotics to derive its optimality in situations where the quality varies, from the numerics we can see that this assumption is not necessary for the system to have strong performance in more realistic, non-asymptotic settings. We can also easily see from the numerics that the simple average can not handle such changes in the underlying quality. After the switch the simple average, denoted by the blue line, is perpetually weighed down by the sequentially stale older reviews making its ability to learn the new true mean rating tortuously slow. Further, in the right panel, we see that both the α-moving average and L-sliding window systems can easily handle multiple changes, just as long as the time between changes is not too short.

To study the precise differences between the three rating systems we isolate their performances in Section 4.2.4. From Section 4.2.4 we can see that both the L-sliding window and α-moving average rating systems quickly converge to the new true service quality after a switch.

Figure 13: Average Rating When Service Quality Changes.

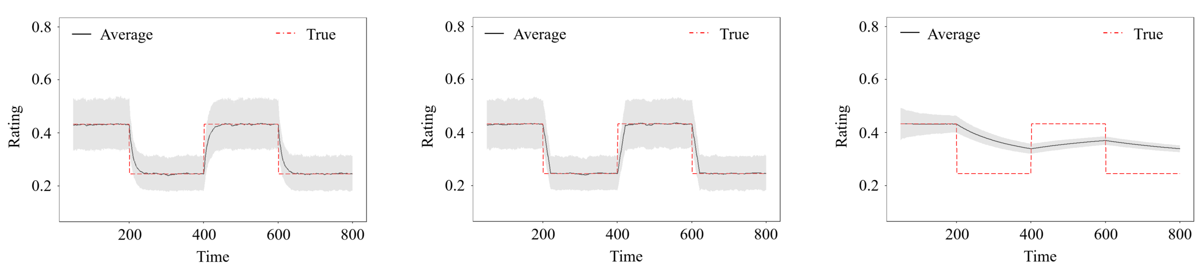

Note. Here we plot the average rating (solid line) v.s. the underlying quality (dashed line) where the ratings for the service provider at each time follow a truncated exponential distribution TExp(λ) and the service quality changes from $\lambda = 4$ to $\lambda = 2$ periodically with time length 200. In the left panel, we plot the average rating of 1000 trials under the α-moving average rating system with $\alpha = 2/21$. In the middle panel, we plot the average rating of 1000 trials under the L-sliding window rating system with $L = 20$. In the right panel, we plot the average rating of 1000 trials under the simple average rating system. The shaded areas are the 90% confidence interval for the average rating.

For Section 4.2.4 a few differences between the two systems become apparent. From the left panel, we can observe that the α-moving average rating system error very quickly drops immediately after the service quality switch occurs, suggesting that it more quickly learns the new rating distribution. However its error comes with a long-tail effect compared to the L-sliding window ratings, which exhibits a slower linear decrease in error at first, but after this learning phase returns to its baseline. Thus, an interesting dichotomy emerges where an α-moving average system may be preferred by platform designers interested in mitigating the worst-case experiences incurred by sequential staling, whereas the L-sliding window may provide a superior mode experience.

In this section we showed that moving-average ratings can effectively mitigate many of the issues due to sequentially staling. However one inherent limitation is of moving average systems, and indeed any rating system, is that they need to be continuously fed new reviews in order to learn about changes in the underlying service quality distribution In the next section we will turn our attention to the temporal staling that occurs when no reviews are written, and describe how a rating system can be extended to prevent these types of gaps in information and incentivize providers to continually solicit new reviews.

4.3 Temporal Staling and Penalty Reviews

In this section, we turn our attention to temporal staling and extending the rating system to include penalty reviews for service providers with long periods of inactivity. Recall that temporal staling is when *no review* has been given for a service provider for a significant period of time. This can occur when a service provider is insufficiently motivated to solicit new reviews, or if they are taking their transactions off platform in order to avoid the platform fee, or if they are no longer offering service and have not notified the hosting platform (see Appendix C.3.1). Temporally staling, much like sequential staling, can cause ratings to not represent the service providers current quality and thus are undesirable. To capture the phenomena of temporal staling we will slightly extend the model introduced in Section 4.2 to now allow for a *no-review* to be given in each period.

If a no-review occurs in some period, say i, then to form the rating we suppose the platform chooses some penalty parameter $\beta \geq 0$ such that the review score in period i is $s_i = \beta$. In this way, if a service provider has a rating that would not be improved by a review score of β, their rating is down-weighted to account for the temporal staling which resulted from the absence of a review in that period. As such we will refer to β as the *penalty review*. We emphasize that any rating system previously studied (simple average, moving average, sliding window, etc.) can all suffer temporal staling without some sort of temporally-based penalization. In this subsection we study how the introduction of penalty reviews impacts steady outcomes for service providers.

Specifically, to study the effects of penalty reviews, we consider a general model of two-sided market dynamics. On one side of the market is a pool of customers who hire service providers, taking into account the potential providers' current rating. On the other side of the market are the providers themselves, who are of heterogeneous quality, where quality represents the probability they provide adequate (i.e. highly rated) service. The platform coordinates interactions between the two sides, taking a percentage fee from each successful transaction, as well as managing and updating the ratings for each provider. Additionally, in our model we allow service providers to initiate disintermediation, allowing them to transact off platform without incurring the platform's fee. We assume a review can only be posted after a transaction occurs on the platform otherwise the customer out of submitting a review for the provider, as is common in practice for platforms like Uber and Airbnb. As no review can be given, the disintermediating service provider always incurs the penalty review in the period. For a depiction of disintermediation in our model see Fig. 14.

In this model, we will be interested in the asymptotic performance of the platform as a function of the penalty review term and the percentage fee chosen by the platform. In the next subsection, we formally describe the model and give some preliminary results about the steady-state outcomes.

Figure 14: Market Disintermediation on Online Platforms.

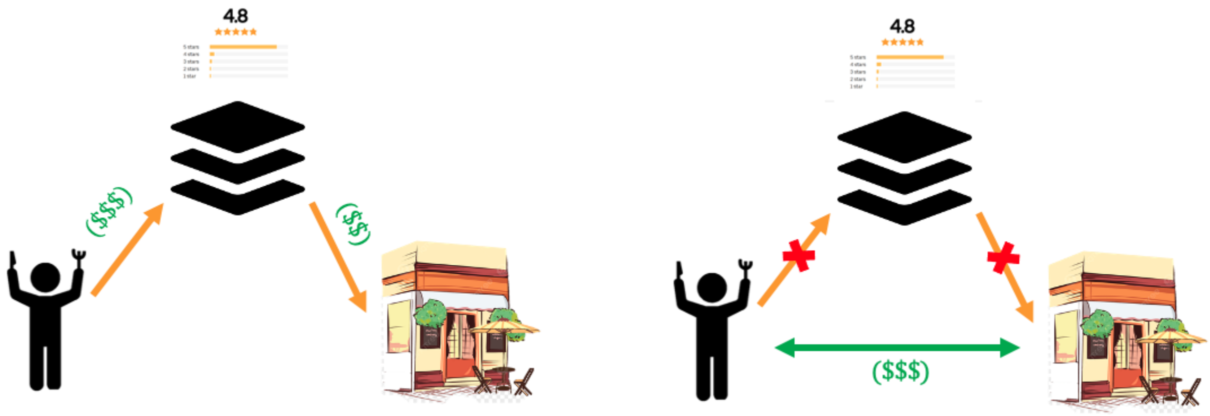

Note. Depicted is an example of market disintermediation with three entities, a customer, a platform, and a service provider (here a restaurant). In the left figure, the customer uses the platform which maintains ratings to discover restaurants. If the customer transacts through the platform, the payment and reviews are handled by the platform for a fee. In the right figure, disintermediation occurs. In this case, the customer discovers the restaurant through the platform but then leaves and transacts directly with the service provider, avoiding the platforms fee but also forgoing the ability to leave a review.

4.3.1 Platform Model and Preliminaries

We suppose each service provider on the platform is described by the pair (F_q, p). F_q is the service providers inherent distribution of service **q**uality, and where $\mathbb{E}[F_q] = q \in (0, 1)$. In fact we will only be concerned with the providers expected quality, however it is notationally useful to be able to refer to the distribution of service quality. The second part of the pair, p, is the probability the provider decides to transact on the platform. That is, once hired, the service provider initiates disintermediation with probability $1 - p$ [3]. The platform observes whether or not a provider transacts on the platform, and if they do transact at some time

[3]We model disintermediation as occurring probabilistically. As we will see, an appropriate choice of penalty β and fee f forces p to the extremes where it is deterministic.

t, observes the subsequent customer review $s_t \sim F_q$. For simplicity we will assume every transaction on the platform results in a review, and the new review is independent of previous ratings. If they do not transact on the platform in the period, the provider incurs a penalty review $s_t = \beta$ where $\beta \in [0, 1]$ which is a priori fixed and chosen by the platform.

The platform maintains an α-moving-average rating system such that a time t the rating for a service provider is $R_t = (1 - \alpha)R_{t-1} + \alpha s_{t-1}$. In each period, and based on their current rating, each service provider has some probability of being hired, $h(R_t) \in (0, 1]$, where $h(\cdot)$ is a link function that maps the providers current rating to the probability of being hired in a period. In each period there are two ways a service provider could fail to receive a review: 1) they could not be hired by any customer which occurs with probability $1 - h(R_t)$, or 2) they are hired but choose to disintermediate which occurs with probability $h(R_t)(1 - p)$. Naturally we assume the platform can not distinguish between a provider not being contracted and a provider being contracted but choosing not to transact on the platform. Finally, the platform decides a fee percentage, f. If the service provider conducts a transaction of value r on the platform, the platform earns fr and the provider receives $(1 - f)r$. If the provider disintermediates, they keep all of r and the platform gets nothing.

The objective of the model is for the platform to choose the penalty term β and fee percentage f such that it maximizes their long run (i.e. asymptotic) per-period expected revenue. To study this objective, fix the moving average rating system parameter α, the transaction revenue r, the link function h, and suppose a service provider earns ratings distributed according to F_q with mean q and, if hired, chooses to transact on the platform with probability p. Given a rating of R_t at time t, the providers expected rating at time $t + 1$ is:

$$\mathbb{E}[R_{t+1}] = (1 - \alpha)R_t + \alpha \left((1 - ph(R_t)) \beta + pq \right), \tag{17}$$

since with probability $ph(R_t)$ the service provider is hired and transacts on the platform incurring an expected rating of $\mathbb{E}[F_q] = q$, and with probability $h(R_t)(1-p) + (1-h(R_t)) = 1 - ph(R_t)$ they are absent in the period and receive the penalty review.

This set-up induces a Stackelberg game in which first the platform announces the penalty β and fee f. Then the service providers, based on their own distributions of service quality F_q, choose their probability of transacting on the platform, p. Asymptotically, this induces some equilibrium outcomes, which under the assumption that the link function h is linear, we can characterize.

Proposition 1 (Market Equilibrium Under α-Moving Average Ratings). *If h is a linear function of the form $h(x) = a + (b-a)x$ for constants $0 \leq a < b \leq 1$, and the platform commits to a penalty term β and fee f, then:*

(a) A provider will not disintermediate (choose $p = 1$) on the platform if their expected service quality q is such that $q \geq \beta + f/(b-a)$, else they will always disintermediate (choose $p = 0$).

(b) For a provider with $q \geq \beta + f/(b-a)$, the platform earns an asymptotic expected per-period profit of $rf\left(\frac{a+\beta(b-a)}{1-(q-\beta)(b-a)}\right)$.

(c) For a provider with $q \geq \beta + f/(b-a)$, the service provider earns an expected per-period profit of $r(1-f)\left(\frac{a+\beta(b-a)}{1-(q-\beta)(b-a)}\right)$.

Proposition 1 characterizes the equilibrium behavior for each service provider in the market given fixed β and f. For providers who can provide sufficiently high quality service, the penalty review disincentives temporal staling by motivating the service provider to always transact on platform and acquire new reviews. For lower quality service providers the opposite is true, and they have no incentive to transact on-platform. Thus one interesting feature of our model revealed by Proposition 1 is that by setting the penalty term and fee, the platform effectively decides a minimum acceptable level of service quality which is equal to $\beta + f/(b-a)$. For any service provider whose inherent expected service quality is less than this floor, they (asymptotically) always have an incentive to transact off the platform and their rating will shrink to a minimum rating of β. For providers with inherent quality in excess of $\beta + f/(b-a)$, the hit to their rating that they incur by transacting off the platform hurts their long-term earning ability (since higher ratings lead to more business in our model) more than it helps by avoiding the fee, motivating them never to disintermediate. Note that a minimum rating for service providers is a common feature of modern two-sided platforms. For instance example, Uber removes drivers with ratings below 4.6 [103]. Proposition 1 gives motivation for these types of market rules, and usefully leverages it to stem profit loss due to disintermediation.

In Proposition 1 we assume h, the function that links a service provider's rating to their probability of being hired in a period, is a linear function. We believe this assumption is relatively mild. Any platform can essentially assign a service provider some probability of being hired simply by varying the providers' placement in the platform display. That is, if a provider is currently receiving a job 80% of the time, that proportion can be reduced in expectation by any fraction $\gamma \in (0, 1)$ simply by omitting the provider from γ proportion of the searches. Similarly, a platform can increase a provider's hiring rate by varying a provider's position in its recommendation algorithm, for instance placing them higher up in the search results. In this way, h is essentially platform controlled, and so for simplicity we will suppose it is implemented as linear in the rating. Note that in the linear implementation, since the rating is always between 0 and 1, the parameters a and b determine the minimum (a) and maximum (b) probability of being hired in a period.

In the remainder of this chapter, we will take this minimum rating implied by Proposition 1 as an exogenously determined parameter, $\tau := \beta + f/(b - a)$, and derive guidelines for how to set the remaining parameters of this system.

4.4 Prior-Free Fresh Rating System Optimization

In this section, we consider how a platform should set the penalty term β and fee f to prevent temporal staling and maximize platform revenue when given distributional information about their service provider's expected quality. Suppose q, the expected service quality for a service provider, is drawn from a distribution G with density g which describes the pool of service providers on the platform. In this stochastic set-up we search for choices of β and f which (approximately) maximize the platform's long-term revenue.

By Proposition 1 we know in limit that any provider with an incentive to disintermediate will always disintermediate, dropping the revenue the platform earns from them to zero. Thus the inequality $q \geq \beta + f/(b - a)$ where a and b are parameters of the linear link function h is a necessary condition for profit; the platform can only extract fees from providers with expected quality in excess of this bound. Since this bound acts as a minimum level quality allowed on the platform, we will treat the desired value of this bound, τ, as exogenously set by the platform. That is we suppose the platform fixes some $\tau \in (0,1)$, and then must choose β and f such that they satisfy $\beta + f/(b-a) = \tau$. We will return to the question of how to choose this τ at the end of this section.

Now, when τ is fixed, $q \sim G$, and h is linear, the platform must solve the following optimization problem to maximize their asymptotic per-period profit:

$$\text{(Platform Profit)} \quad \max_{\substack{\beta+f/(b-a)=c, \\ \beta, f\in[0,1]}} r \int_c^1 f\left(\frac{a+\beta(b-a)}{1-(q-\beta)(b-a)}\right) g(q)dq \tag{18}$$

When the distribution of expected service quality G is known precisely, Eq. (18) can be solved numerically. However G is typically not fully known to the platform. Moreover, even if the platform could estimate G, there is no reason to believe this distribution will remain constant over time as service providers come and go on the platform, and as market conditions shift. Thus, instead of trying to tailor the penalty and fee specifically to the distribution of expected service quality, we will aim to find robust specifications of the parameters (β, f) that induce approximately optimal revenue for *all* possible service distributions. As such a specification does not depend on the exact form of the service distribution (i.e. the prior) such a guarantee is referred to as *prior-free*. Remarkably, such a specification of parameters is indeed possible.

Theorem 12 (Approximate Rating Design)**.** *For any level of minimum quality τ, any linear link function h, and any distribution $q \sim G$ supported on $[\tau, 1]$, the choice of parameters:*

$$(\beta, f) = \left(h^{-1}\left(\frac{\left(\sqrt{1-h(\tau)}\right)h(\tau)}{2\left(1+\sqrt{1-h(\tau)}\right)}\right), \frac{\left(2+\sqrt{1-h(\tau)}\right)h(\tau)}{2\left(1+\sqrt{1-h(\tau)}\right)}\right)$$

gives a $\frac{2+\sqrt{1-h(\tau)}}{2\left(1+\sqrt{1-h(\tau)}\right)} \geq 3/4$ *approximation of the optimal achievable expected platform revenue.*

Theorem 12 yields approximately optimal choices for the penalty and fee parameters as a function of τ, the minimum acceptable quality of service on the system, and h the link function. It states for any choice of τ, the above parameters are such that $\tau = \beta + f/(b-a)$ ensuring no incentive to disintermediate, and they guarantee at least $3/4^{th}$ of the maximum achieveable revenue. This is relative to the optimal parameter choices which could depend on the distribution of expected service quality among participants, G, in a detailed way. Further note our approximation guarantee is often much stronger than $3/4^{th}$ which is obtained only in the case that $h(\tau)$ tends to 0. For intermediate choices of τ the guarantee quickly improves, see the right panel Fig. 16 for a visualization of the bound. Moreover, for every choice of τ the guarantee above is tight in the sense that, when h is the identity function, there is a simple distribution such that the above prior-free choices of (β, f) give only $\frac{1-\sqrt{1-\tau}+\tau}{2\tau}$ of the optimal achieveable revenue. For details of this construction, see Example 9. Finally note the choice of β above has been written in terms of the inverse function of h. Since h is a linear function, its inverse is $h^{-1}(x) = \frac{x-a}{b-a}$.

For platform designers, Theorem 12 gives easily implementable guidance for how to set the penalty review parameter and the associated fee. Once the platform sets a minimum service level τ, they can then plug in the heuristic parameter values to garner almost all the achieveable revenue. However, to derive our prior-free approximately optimal choices of β and f, we required τ to be exogenously set by the platform. A natural question then how to set τ? In general, we believe the nature of online rating distributions often implies a natural choice of τ. We explore this idea in the following remark, and numerically in the next subsection.

Remark 2 (Choosing τ)**.** *Many platforms set a minimum rating threshold on their platform and enforce this threshold by removing service providers who drop below it. In our model we call this threshold τ, and here consider how a platform might choose it. In general, note that the dependence between τ and platform profit expressed in Eq.* (18) *is quite complicated. While we can find good prior-free choices for the penalty and fee given τ, jointly choosing all three in a way which provably yields the majority of the revenue is more challenging. Intuitively, a good of choice of τ is one that enforces a high level of minimum quality while also removing only a small fraction of service providers. Thankfully, in practice the distribution* *of server ratings often takes a simple form which yields a natural choice of τ with these intuitive properties.*

[138] shows that distribution of ratings on Yelp.com are unimodal and left-skewed, and this shape for the rating distribution on Yelp has remained consistent for as long as the company has released data[4] *For an example of such distributions, see the left panels of Section 4.4.1 and Fig. 16. By assuming the distribution of expected ratings is left-skewed and unimodal, we can then reason that a good choice for τ should be one on the left tail since the majority of ratings are relatively good and clustered around 4 stars. Thus we propose choosing τ in accordance with the Pareto principle [47]. That is, to set it at the 20^{th} percentile so that $\tau = G^{-1}(.2)$. This is in line with standard Pareto principle reasoning that 80% of an effect stems from 20% of the causes. Here the effect is negative reviews and the cause are the bottom 20% of service providers. Finally, we note that this choice of τ is not precisely prior-free as it depends on G, however the general principle that τ should be set conservatively to remove only the worst service providers is not distribution specific and should be implementable for a platform. We numerically study the performance of choosing τ at the 20^{th} percentile for a class of left-skewed unimodal distributions in Fig. 16.*

4.4.1 Penalty Review Numerics

In this subsection, we will numerically explore the effectiveness of the heuristic policy implied by Theorem 12 and Remark 2. All code for these numerics can be found at `https://github.com/tcui-pitt/Ratings`.

In our first set of numerics in Section 4.4.1, we consider the setting where the service providers' service quality follows the rating distribution from Yelp.com [138] , which is discretized into five buckets corresponds to 1-5 star ratings. See the left panel of Section 4.4.1 for the associated density function. For this distribution, we suppose different levels of τ are exogenously chosen by the platform and examine how closely the heuristic choice of parameters given by Theorem 12 approximates the optimal distribution-specific choices.

[4]See `https://www.yelp.com/dataset` for this data.

Figure 15: Optimal Expected Revenue vs. Heuristic Revenue as τ Varies

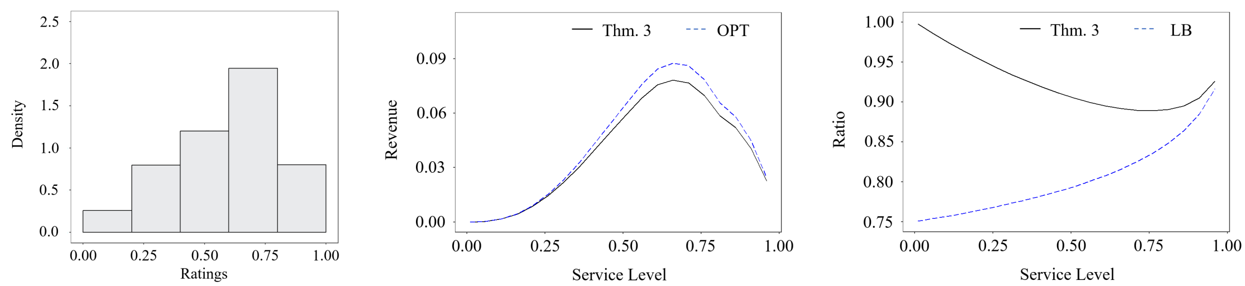

Note. We visualize the difference between the revenue from the optimal distribution-specific penalty and fee (β^*, f^*) versus the revenue of heuristic choice of (β, f) as in Theorem 12 as τ varies for an archetypal left-skewed unimodal distribution. In the left panel, we plot the histogram of expected service quality as reported by [138] where each bucket corresponds to the number of stars. In the middle panel, we plot the heuristic revenue versus the optimal revenue as τ varies. In the right panel, we plot the ratio of revenue of the heuristic policy vs the optimal revenue and the lower bounding ratio from Theorem 12.

In the middle panel of Section 4.4.1, we plot expected revenue when using the heuristic fee f and penalty β in Theorem 12 for any given service level τ. We find that the heuristic very closely approximates the true optimal policy when τ is specified exogenously. Indeed in the right panel of Section 4.4.1 we see that these choices of parameters greatly out-perform the worst case lower bound garner in excess of 90% of the optimal revenue for almost all τ.

Next, we consider the performance of the 20^{th} percentile rule described in Remark 2. Specifically, we consider a class of left-skewed unimodal distributions defined by $q \sim \text{Beta}(x, 2)$ where x varies between 5 and 15. For each distribution in this class, we heuristically choose τ, f and β and compare it against the revenue maximizing choices of these parameters (computed via grid search).

Figure 16: Optimal Service Level vs. Heuristic Service Level.

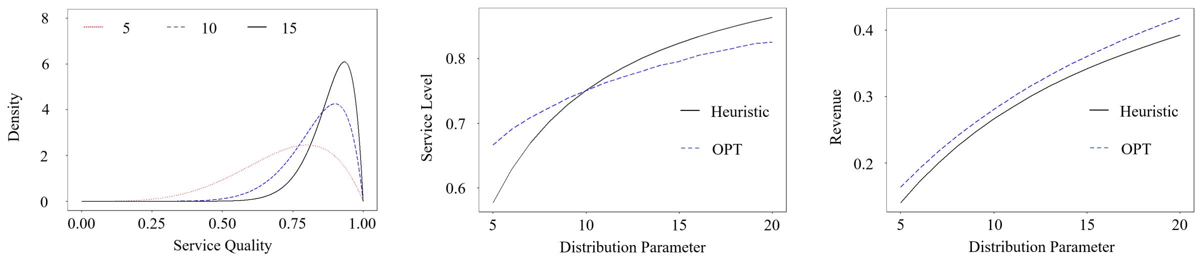

Note. We visualize the performance of the heuristic policy where τ is chosen as as 20% quantile of the Beta(x, 2) distributions, and (β, f) are chosen as in Theorem 12, as x varies. In the left panel, we plot three Beta(x, 2) distributions with parameters $x = 5, 10, 15$, respectively. In the middle panel, we plot the optimal service level and the heuristic service level (20% quantile of the Beta(x, 2) distribution) when the distribution parameter varies from $x = 5$ to $x = 20$. In the right panel, we plot the optimal revenue and the revenue generated by our heuristic service level with heuristic choices of for (β, f).

In the left panel of Fig. 16, we plot the density of the Beta(x, 2) distributions for $x = 5, 10, 15$ respectively, from which we see that the Beta(10, 2) distribution approximates the distribution of average business ratings in [138]. In the middle panel of Fig. 16, we plot the heuristic choice of the service level τ where the τ is chosen as the 20^{th} percentile of the distributions, and the optimal service level which is found by grid search. From Fig. 16, we see when x is large (for example $x > 10$), the distribution is concentrated around the highest levels of service quality, and the heuristic choice of service level τ is close to the optimal service level, suggesting the Pareto principle may lead to effective choices of τ in practice. Finally, in the right panel we plot the expected revenue generated using the heuristic choice of τ (also f and β as specified in Theorem 12) and the optimal service level (also optimal f and β) found by grid search. We observe the expected revenue of the heuristic choice is consistently close to the optimal expected revenue.

Thus, the numerical results in Fig. 16 imply the heuristic choices of τ, β and f describe above should work well in practice.

4.5 Conclusions and Future Directions

In this work, we considered the problem of staling in online rating systems. Staling occurs when old, irrelevant reviews inform current ratings rendering the ratings unreliable for customers. Many simple rating system implementations are plagued by issues due to staling, chief among them being the simple average which, despite its wide use, can easily become stale. We made a distinction between two types of staling, sequential staling and temporal staling, each of which required different interventions to mitigate. For sequential staling we introduced and studied the α-moving average rating system which places proportional weight on the latest review. We showed that among all oblivious rating systems, the moving average rating system provides the best protection against staling by exhibiting min-max mean squared error when a change in the underlying quality could occur. For temporal staling, we introduced the concept of penalty reviews which down-weight provider ratings when there are no new reviews. In a general model for two-sided platforms, we connected the notion of temporal staling with market disintermediation and derived new theory to guide how a platform should choose the penalty review term and fee parameters. Specifically, we gave a prior-free heuristic specification of the parameters that provably encourages service providers to solicit new reviews (by transacting on the platform) and (approximately) maximizes profit.

Taken together, our work can be seen as advocating for a suite of tools that platforms can turn to when issues regarding staling occur. Our theory gives support for the use of moving average rating systems in practice. In many ways, moving average systems can be seen as an evolution of L-sliding window policies. Moving average rating systems offer a smooth version of L-sliding windows, inheriting the same virtues of L-sliding window systems while also being more responsive to changes in service provider quality. α-moving average ratings are also simple to implement and understand, they are always correct in expectation when the distribution of service quality is stationary, and they can be set to any arbitrary level of consistency (i.e. variance in the rating error) simply by tuning its one parameter, α.

We also propose the use of penalty reviews for service providers who have not received recent reviews. There are many reasons why a service provider's rating may become temporally stale, in our model we focused on the specific case of temporal staling due to market disintermediation. For this important use case, penalty reviews can completely disincentive high quality providers from transacting off-platform while pruning lowquality and inactive providers from the market. While we emphasized the lens of disintermediation, in general, penalty reviews always provide incentives for service providers on the platform to continually solicit new reviews. We believe this sustained pressure can improve overall rating consistency in outcomes for users, and encourage service providers to always provide the highest quality service possible. Note that these penalty reviews are simple to implement, intuitive, and easy to optimize since they consist of just one tunable parameter (β, the penalty term) which can be put into correspondence with a minimum rating for providers.

Finally, we emphasize the necessity of our proposed rating modifications. The world is in constant flux and change is inevitable. No service level can be held constant forever. Thus any rating system that does not down-weight older reviews will eventually be dominated by those older reviews and made irrelevant. Any rating system that does not monitor how long it has been since the last review will eventually find itself frozen in time. Our results study simple, natural ways for rating systems to adapt to these eventualities.

We believe our work opens up many avenues for future work relating to issues of staling. Note that even though ratings are a deeply ingrained part of the way a typical customer conducts commerce, there is no standard way to build rating systems. Google and Yelp use a simple average, Uber uses an L-sliding window system, and many other companies are currently experimenting with the structure of their rating system. For instance, Amazon.com uses a proprietary rating system that does not seem to follow any of the formats listed above [95]. There is a noticeable dearth of empirical work regarding principled ways that current platforms can deal with issues due to staling. In future work, we believe it would be productive to understand the interplay between our theoretically motivated anti-staling interventions, and some of the specific use cases and constraints encountered in practice.

On the theoretical side, in this work, we emphasize interventions that are static; they are characterized by single fixed parameters. It would be interesting to look at dynamic ways to set α and β. For instance, in the related problem of change point detection (see the literature review for relevant papers), the dominant modern approach is to use statistical testing to determine if a change in the underlying service distribution has occurred. It would be interesting if these ideas could be incorporated into a dynamic choice of α that grows when the detected risk of a change is high. Similarly, for the penalty term β, this penalty is applied regardless of the providers' current rating. It would be interesting to study the power of penalty terms that are specific to providers' current rating, and perhaps also include momentum-based terms to detect when providers have become inactive. We hope to study some of these directions in future work.

Appendix A Feature-Based Pricing

A.1 Omitted Examples

Example 1 ($p_\epsilon(x)$ can be discontinuous)**.** *Suppose* $\mu(\mathbb{X}) \sim$ *Uniform*$[1,2]$ *(or any other continuous distribution on [1,2]) and* ϵ *is either -.5 or .5 with probability* $\frac{1}{2}$*. Then for every* $x \leq 1.5$*, the optimal price is* $p_\epsilon(x) = x + .5$*, and* $p_\epsilon(x) = x - .5$ *otherwise. Thus at 1.5* $p_\epsilon(x)$ *is discontinuous, and by Lemma 1 the revenue function* $\mathcal{R}_\epsilon(x)$ *is non-differentiable.*

Example 2 (Optimal Segmentations Need Not be Interval)**.** *In this example we give* $\mu(\mathbb{X})$ *and* ϵ *such that the optimal segmentation and pricing is non-interval. Specifically, for any number* $k \geq 2$*, assume* $\mu(\mathbb{X})$ *is uniformly distributed on the set* $\{k, k+2, k+4, \cdots, 3k\}$*, and* ϵ *is either* k *or* $-k$ *with probability* $\frac{1}{2}$*, respectively. Note that if we consider fully personalized pricing, the optimal price for* $\{k\}$ *is* $2k$*, and the optimal price for* $\{3k\}$ *is either* $2k$ *or* $4k$*. Therefore, the unique optimal k-market segmentation and pricing uses segments,*

$$\mu(\mathcal{X}_1) = \{k, 3k\}, \mu(\mathcal{X}_2) = \{k+2\}, \cdots \mu(\mathcal{X}_k) = \{k + 2(k-1)\},$$

with corresponding price for each segment,

$$p(\mathcal{X}_1) = 2k, p(\mathcal{X}_2) = 2k+2, \cdots, p(\mathcal{X}_k) = 2k + 2(k-1).$$

To see this segmentation achieves the optimal revenue note the revenue it achieves is the same as from fully personalized pricing. Further since $p_\epsilon(k) = p_\epsilon(3k)$*, and this is not true for any other predicted valuations, no other* k*-segmentation can achieve the same revenue. As the first segment is not interval, the optimal segmentation thus needs not to be interval for any* k*.*

Example 3 (Tightness of Theorem 4)**.** *Suppose the regression model has no error, i.e. $V|\mathbf{x} = \mu(\mathbf{x})$, and let $\mu(\mathbb{X}) \sim$ Uniform$[0,t]$ for some $t > 0$. Then, $\mathbb{E}[\mathcal{R}_\epsilon(\mu(\mathbb{X}))] = \mathbb{E}[V] = \frac{t}{2}$. To compute $\mathcal{R}_{kXP}$ for some k, note by Theorem 2 the optimal segmentation here is interval and further each segment can be described by a left and right endpoint in the space of predicted valuations. Let $0 < s_0 < \ldots < s_k = t$ describe those segments (i.e. $\mathcal{X}_i = \{x|\mu(x) \in [s_{i-1}, s_i)\}$) with corresponding prices $p_1, \ldots, p_k$. It is easy to see since $\epsilon = 0$, the optimal price and segmentations must satisfy $s_{i-1} = p_i$ for $i = 1, \ldots, k$ since, if not, increasing the segment interval s_{i-1} up to p_i only increases revenue.*

Now, on segment $\mathcal{X}_i$, the conditional distribution of V is still uniform, so the contribution of that segment to $\mathbb{E}[V]$ is $\frac{s_i+s_{i-1}}{2} \cdot \frac{s_i-s_{i-1}}{t}$ for all i, since only $\frac{s_i-s_{i-1}}{t}$ fraction of the market is in this interval. By contrast, for $i = 1, \ldots, k$, the k-market segmentation strategy on segment i earns revenue $p(\mathcal{X}_i)\Pr(\mu(\mathbb{X}) \geq p(\mathcal{X}_i)|\mathcal{X} \in [s_{i-1}, s_i)\Pr(\mathbb{X} \in [s_{i-1}, s_i) = s_{i-1}\frac{s_i-s_{i-1}}{t}$ since $p_i = s_{i-1}$, and thus all customers in the segment buy. The difference in revenue is then

$$\begin{aligned}
\mathbb{E}[\mathcal{R}_\epsilon(\mu(\mathbb{X}))] - \mathcal{R}_{kXP} &= \frac{s_0^2}{2t} + \sum_{i=1}^{k} \frac{s_i + s_{i-1}}{2} \cdot \frac{s_i - s_{i-1}}{t} - s_{i-1}\frac{s_i - s_{i-1}}{t} \\
&= \frac{s_0^2}{2t} + \frac{1}{2t}\sum_{i=1}^{k}(s_i - s_{i-1})^2 = \frac{1}{2t}\left((s_0 - 0)^2 + \sum_{i=1}^{k}(s_i - s_{i-1})^2\right).
\end{aligned}$$

By inspection, for a fixed $s_0 = \frac{t}{k+1}$, the segmentation which minimizes this difference is equispaced, i.e., $s_i = s_{i-1} + \frac{t}{k+1}$ for $i = 1, \ldots, k$. Plugging in gives $\mathbb{E}[\mathcal{R}_\epsilon(\mu(\mathbb{X}))] - \mathcal{R}_{kXP} = \frac{t}{2(k+1)} = \frac{E[V]}{k+1}$.

A.2 Omitted Proofs

A.2.1 Omitted Proofs from Section 2.2

Proof of Lemma 1. **(a)** Fix some ϵ and positive real numbers x_1, x_2 such that $x_1 < x_2$ and recall $\theta_\epsilon(x) := p_\epsilon(x) - x$ is the difference between the price and x. Further recall $p_\epsilon(x_1), p_\epsilon(x_2)$

are prices that maximize $p\overline{\mathrm{F}}_\epsilon(p - x_1)$ and $p\overline{\mathrm{F}}_\epsilon(p - x_2)$ respectively. Thus, by optimality we have the following two inequalities

$$(x_1 + \theta_\epsilon(x_1))\, \overline{\mathrm{F}}_\epsilon(\theta_\epsilon(x_1)) \geq (x_1 + \theta_\epsilon(x_2))\, \overline{\mathrm{F}}_\epsilon(\theta_\epsilon(x_2)), \tag{19}$$

$$(x_2 + \theta_\epsilon(x_2))\, \overline{\mathrm{F}}_\epsilon(\theta_\epsilon(x_2)) \geq (x_2 + \theta_\epsilon(x_1))\, \overline{\mathrm{F}}_\epsilon(\theta_\epsilon(x_1)). \tag{20}$$

Rearranging the two inequalities yields,

$$\frac{x_1 + \theta_\epsilon(x_1)}{x_1 + \theta_\epsilon(x_2)} \geq \frac{\overline{\mathrm{F}}_\epsilon\left(\theta_\epsilon(x_2)\right)}{\overline{\mathrm{F}}_\epsilon\left(\theta_\epsilon(x_1)\right)} \geq \frac{x_2 + \theta_\epsilon(x_1)}{x_2 + \theta_\epsilon(x_2)}.$$

Consequently,

$$\left(x_1 + \theta_\epsilon(x_1)\right)\left(x_2 + \theta_\epsilon(x_2)\right) \geq \left(x_1 + \theta_\epsilon(x_2)\right)\left(x_2 + \theta_\epsilon(x_1)\right),$$

Simplifying the expression, we get

$$(x_2 - x_1)\, \theta_\epsilon(x_1) \geq (x_2 - x_1)\, \theta_\epsilon(x_2).$$

Finally, noting $x_2 - x_1 > 0$, the inequality is equivalent to $\theta_\epsilon(x_1) \geq \theta_\epsilon(x_2)$ and thus the margin monotone decreasing.

(b) As in (a), fix some ϵ and positive real numbers x_1, x_2 such that $x_1 \leq x_2$. Then $x_1 + \epsilon \leq_{\mathrm{st}} x_2 + \epsilon$ in the sense of first order stochastic dominance, and it is well known that stochastic dominance of the valuations implies $\mathcal{R}_\epsilon(x_1) \leq \mathcal{R}_\epsilon(x_2)$ (see for instance [65] for an extended discussion). Combining this observation with Eqs. (19) and (20) above yields,

$$\mathcal{R}_\epsilon(x_2) - \mathcal{R}_\epsilon(x_1) \geq (x_2 + \theta_\epsilon(x_1))\, \overline{\mathrm{F}}_\epsilon(\theta_\epsilon(x_1)) - (x_1 + \theta_\epsilon(x_1))\, \overline{\mathrm{F}}_\epsilon(\theta_\epsilon(x_1)) = (x_2 - x_1)\overline{\mathrm{F}}_\epsilon(\theta_\epsilon(x_1)),$$

$$\mathcal{R}_\epsilon(x_2) - \mathcal{R}_\epsilon(x_1) \leq (x_2 + \theta_\epsilon(x_2))\, \overline{\mathrm{F}}_\epsilon(\theta_\epsilon(x_2)) - (x_1 + \theta_\epsilon(x_2))\, \overline{\mathrm{F}}_\epsilon(\theta_\epsilon(x_2)) = (x_2 - x_1)\overline{\mathrm{F}}_\epsilon(\theta_\epsilon(x_2)).$$

Dividing both sides by $x_2 - x_1$ gives,

$$\overline{\mathrm{F}}_\epsilon(\theta_\epsilon(x_1)) \leq \frac{\mathcal{R}_\epsilon(x_2) - \mathcal{R}_\epsilon(x_1)}{x_2 - x_1} \leq \overline{\mathrm{F}}_\epsilon(\theta_\epsilon(x_2)). \tag{21}$$

When p_ϵ is continuous then θ_ϵ is also continuous, and taking $x_1 \to x_2$ squeezes the derivative to be $\overline{\mathrm{F}}_\epsilon(p_\epsilon(x) - x)$ as desired.

(c) $\mathcal{R}_\epsilon(x)$ was noted to be increasing in the proof of (b). Now to prove continuity, fix ϵ and positive real numbers x_1, x_2 such that $x_1 < x_2$. Then,

$$\mathcal{R}_\epsilon(x_1) \leq \mathcal{R}_\epsilon(x_2) = (x_2 + \theta_\epsilon(x_2) + (x_2 - x_1) - (x_2 - x_1))\,\overline{\mathrm{F}}_\epsilon(\theta_\epsilon(x_2)) \leq \mathcal{R}_\epsilon(x_1) + (x_2 - x_1),$$

where the last inequality follows from distributing and applying Eq. (20), and the fact that $\overline{\mathrm{F}}_\epsilon(\cdot) \leq 1$. Taking $x_1 \to x_2$ gives us the continuity of $\mathcal{R}_\epsilon(x)$.

For convexity, again fix positive real numbers x_1, x_2 and also $\lambda \in (0, 1)$. Then,

$$\begin{aligned}\mathcal{R}_\epsilon(x_1) &= p_\epsilon(x_1)\overline{\mathrm{F}}_\epsilon(p_\epsilon(x_1) - x_1)\\ &\geq (p_\epsilon(\lambda x_1 + (1-\lambda)x_2) + (1-\lambda)(x_1 - x_2))\,\overline{\mathrm{F}}_\epsilon(p_\epsilon(\lambda x_1 + (1-\lambda)x_2) - (\lambda x_1 + (1-\lambda)x_2)),\end{aligned}$$

where the inequality follows from noting that $p_\epsilon(x_1)$ is revenue optimal for $x_1 + \epsilon$ and any other price can earn no more. Similarly,

$$\begin{aligned}\mathcal{R}_\epsilon(x_2) &= p_\epsilon(x_2)\overline{\mathrm{F}}_\epsilon(p_\epsilon(x_2) - x_2)\\ &\geq (p_\epsilon(\lambda x_1 + (1-\lambda)x_2) - \lambda(x_1 - x_2))\,\overline{\mathrm{F}}_\epsilon(p_\epsilon(\lambda x_1 + (1-\lambda)x_2) - (\lambda x_1 + (1-\lambda)x_2)).\end{aligned}$$

Combine the two inequalities above, we get

$$\lambda\mathcal{R}_\epsilon(x_1) + (1-\lambda)\mathcal{R}_\epsilon(x_2) \geq \mathcal{R}_\epsilon(\lambda x_1 + (1-\lambda)x_2)),$$

which means $\mathcal{R}_\epsilon(x)$ is convex in x. $\square$

A.2.2 Omitted Proofs from Section 2.3

Proof of Theorem 1. We will prove hardness by showing the Hitting set problem can be reduced to an instance of k feature-based market segmentation and pricing (kXP). Let $\mathcal{X}$ be the ground set of elements of size, $|\mathcal{X}| = m$, and let $\{H_i\}_{i=1}^{n}$ be a collection of subsets of $\mathcal{X}$. Consider the decision version of the hitting set problem, which asks whether there exists a subset of $\mathcal{X}^* \subset \mathcal{X}$, $|\mathcal{X}^*| \leq k$, such that $\mathcal{X}^*$ has non-empty intersection with each H_i. To build a corresponding k-market segmentation and pricing problem, suppose we have n customers such that each customer's valuation is $x_i + \epsilon$ (equivalently, $\mu(\mathbb{X})$ is uniformly supported on these valuations), where $i = 1, 2, ..., n$. Let $p_j = n + \frac{j-1}{m}$, for $j = 1, 2, ..., m$, and let $x_1 = p_1$, and $x_i = x_{i-1} + \frac{p_1}{2(i-1)} + \frac{p_m}{2i}$, for $i = 2, ..., n$. We will now construct an $\epsilon \sim F_\epsilon$ such that, for each x_i, $\mathcal{R}_\epsilon(x_i)$ is maximized at price p_j if and only if in the hitting set problem the subset H_i contains element x_j.

Our construction of ϵ is supported on numbers of the form $p_i - x_j$. Before constructing ϵ, note that p_j is strictly increasing in j, and that $p_j - x_i < p_{j'} - x_i$ as long as $j' > j$. Further note $p_m - x_{i+1} < p_1 - x_i$ since by the definition of x_i and x_{i+1}, $x_{i+1} - x_i > \frac{p_1}{2n} + \frac{p_m}{2n} > 1$ and since $p_m - p_1 = \frac{m-1}{m} < 1$. Let $t_{1,1} = p_1 - x_n$, ..., $t_{j,i} = p_j - x_{n+1-i}$, ..., $t_{m,n} = p_m - x_1$, and let $t_{0,n} = -x_n$, and $t_{m,n+1} = p_m$. Thus, we have

$$t_{0,n} < t_{1,1} < t_{2,1} < \ldots < t_{m,i} < t_{1,i+1} < t_{2,i+1} < \ldots t_{m-1,n} < t_{m,n} < t_{m,n+1}. \tag{22}$$

Now we are ready to define the complementary cumulative distribution function (cCDF) of ϵ. We will let ϵ be such $\overline{\mathrm{F}}_\epsilon(t_{0,n}) = 1$, $\overline{\mathrm{F}}_\epsilon(t_{m,n+1}) = 0$, and working backwards recursively from $\overline{\mathrm{F}}_\epsilon(t_{m,n+1})$ as follows:

$$\overline{\mathrm{F}}_\epsilon(t_{j,i}) = \begin{cases} \frac{i}{p_j}, & \text{if } x_j \in H_i \\ \overline{\mathrm{F}}_\epsilon(t_{j+1,i}), & \text{if } x_j \notin H_i \text{ and } j < m \\ \overline{\mathrm{F}}_\epsilon(t_{1,i+1}) & \text{if } x_j \notin H_i \text{ and } j = m,\ i < n \\ 0 & \text{otherwise.} \end{cases}$$

Note this construction is well defined and is quadratically supported, an example what $\overline{\mathrm{F}}_\epsilon$ looks like is provided in Fig. 22. Further, for any value t such that $t_{j,i} \leq t < t_{j+1,i}$,

$\overline{\mathrm{F}}(t) = \overline{\mathrm{F}}(t_{j,i})$. Now we need to check that $\overline{\mathrm{F}}_\epsilon$ is non-increasing and thus a properly defined cCDF, and also that $p_j\overline{\mathrm{F}}_\epsilon(p_j - x_i)$ is revenue-maximizing only when j, i are such that $x_j \in H_i$.

To the first point, since $p_j \geq n$ for all $j = 1, 2, ..., m$, and $\{p_j\}$ is increasing, therefore, $\frac{i}{p_j} < \frac{i}{p_{j+1}}$. Then, to show $\overline{\mathrm{F}}_\epsilon$ is non-increasing, we only need to show $\frac{i}{p_m} < \frac{i-1}{p_1}$. Note that

$$\begin{aligned}\frac{i}{p_m} - \frac{i-1}{p_1} &= \frac{ip_1 - (i-1)p_m}{p_1 p_m}\\ &= \frac{i(p_1 - p_m) + p_m}{p_1 p_m} > 0,\end{aligned}$$

where the inequality follows from the fact that $p_m - p_1 = \frac{m-1}{m} < 1$ for $1 \leq i \leq n$, and $p_m > n$. Thus $\overline{\mathrm{F}}_\epsilon$ is non-increasing, i.e., $\overline{\mathrm{F}}_\epsilon$ is a proper cumulative distribution function.

Next, we show that $p\overline{\mathrm{F}}_\epsilon(p - x_i) = i$ iff $p = p_j$ and $x_j \in H_i$, and for all other prices the revenue $p\overline{\mathrm{F}}_\epsilon(p - x_i)$ is strictly less than i. By the definition of $\overline{\mathrm{F}}_\epsilon(p_j - x_i)$, if $x_j \in H_i$, $\overline{\mathrm{F}}_\epsilon(p_j - x_i) = \frac{i}{p_j}$, consequently, $p_j\overline{\mathrm{F}}_\epsilon(p_j - x_i) = i$. So now suppose price p satisfies $\overline{\mathrm{F}}_\epsilon(p - x_i) = \overline{\mathrm{F}}_\epsilon(p_{j'} - x_{i'}) = \frac{i'}{p_{j'}}$. To simplify the discussion, we take the largest price p such that $\overline{\mathrm{F}}_\epsilon(p - x_i) = \overline{\mathrm{F}}_\epsilon(p_{j'} - x_{i'})$, i.e., $p - x_i = p_{j'} - x_{i'}$, and by rearranging $p = p_{j'} - x_{i'} + x_i$. All other prices less than $p_{j'} - x_{i'} + x_i$ and which satisfies $\overline{\mathrm{F}}_\epsilon(p - x_i) = \overline{\mathrm{F}}_\epsilon(p_{j'} - x_{i'}) = \frac{i'}{p_{j'}}$ will give us less revenue. Now we want to show that $p\overline{\mathrm{F}}_\epsilon(p - x_i) < i$, i.e.,

$$(p_{j'} - x_{i'} + x_i)\,\frac{i'}{p_{j'}} < i.$$

If $i' < i$, the inequality is the same as

$$x_i - x_{i'} \leq \frac{i - i'}{i'} p_{j'}.$$

By the definition of p_j and x_i,

$$x_{i+1} - x_i = \frac{p_1}{2i} + \frac{p_m}{2(i+1)} < \frac{p_1}{i},$$

where the inequality comes from $\frac{p_m}{p_1} < \frac{n+1}{n}$. Therefore,

$$x_i - x_{i'} \leq \sum_{j=i'}^{i} \frac{p_1}{j} < \frac{i - i'}{i} p_1 < \frac{i - i'}{i} p_j.$$

Similarly, if $i' > i$, $p\overline{\mathrm{F}}_\epsilon(p - x_i) < i$ is equivalent to

$$x_{i'} - x_i > \frac{i' - i}{i'} p_{j'}.$$

Now, by the definition of p_j and x_i,

$$x_{i+1} - x_i = \frac{p_1}{2i} + \frac{p_m}{2(i+1)} > \frac{p_m}{i+1},$$

where the inequality comes from $\frac{p_m}{p_1} < \frac{n+1}{n}$. Therefore,

$$x_{i'} - x_i \geq \sum_{j=i'}^{i} \frac{p_m}{j+1} > \frac{i - i'}{i} p_m > \frac{i - i'}{i} p_j,$$

as desired.

Finally, to determine whether there exists a subset of $X^* \subset X$, $|X^*| \leq k$, such that X^* has non-empty intersection with each H_i, it is equivalent to determine whether there is a k feature-based market segmentation and pricing that yields the maximum revenue $1 + 2 + 3 + \ldots + n = \frac{n(n+1)}{2}$. Since the hitting set problem NP-hard, thus FBMSP is also NP-hard. $\square$

Proof of Lemma 2. **(a)** First we will show $p_\epsilon(x)$ is increasing. Since f_ϵ is log-concave, $\overline{\mathrm{F}}_\epsilon$ is also log-concave (see [14] for an extensive overview of the transformations that preserve log-concavity). Further, $\frac{d}{dx} \log(\overline{\mathrm{F}}_\epsilon(x)) = \frac{-f_\epsilon(x)}{\overline{\mathrm{F}}_\epsilon}$ which by concavity implies the inverse hazard rate, $\frac{\overline{\mathrm{F}}_\epsilon(x)}{f_\epsilon(x)}$, is decreasing in x. Thus $p_\epsilon(\cdot)$ is unique, satisfies first order conditions for revenue optimality, $\frac{d}{dp} p\overline{\mathrm{F}}_\epsilon(p - x)|_{p=p(x)} = 0$, and can be written as $p_\epsilon(x) = \frac{\overline{\mathrm{F}}_\epsilon(p_\epsilon(x)-x)}{f_\epsilon(p_\epsilon(x)-x)}$. Recalling by Lemma 1(a) $p_\epsilon(x) - x$ is decreasing, it thus follows that $p_\epsilon(x)$ must be an increasing function of x.

(b) As in (a) note, if $f(x)$ is log-concave, $\overline{\mathrm{F}}_\epsilon$ is a log-concave function, thus it has Pólya frequency of order 2 (PF2), which is equivalent to that statement that, for any real numbers x_1, x_2, and y_1, y_2, such that $x_1 < x_2$ and $y_1 < y_2$, then $\frac{\overline{\mathrm{F}}_\epsilon(x_1-y_2)}{\overline{\mathrm{F}}_\epsilon(x_1-y_1)} \leq \frac{\overline{\mathrm{F}}_\epsilon(x_2-y_2)}{\overline{\mathrm{F}}_\epsilon(x_2-y_1)}$ (see [112], Section 11).

Now, let $\{\mathcal{X}_i\}_1^k, \{p_\epsilon(\mathcal{X}_i)\}_1^k$ be the optimal segmentation and pricing and suppose WLOG that the prices are distinct. Further suppose the optimal segmentation was not interval, then there exists $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3$ such that $\mu(\mathbf{x}_1) < \mu(\mathbf{x}_2) < \mu(\mathbf{x}_3)$, but with $\mathbf{x}_1, \mathbf{x}_3 \in \mathcal{X}_i$, and $\mathbf{x}_2 \in \mathcal{X}_j$

for some $i \neq j$. Suppose $p_\epsilon(\mathcal{X}_i) < p_\epsilon(\mathcal{X}_j)$ (the opposite case when $p_\epsilon(\mathcal{X}_i) > p_\epsilon(\mathcal{X}_j)$ follows by an identical argument, swapping $\mathbf{x}_3$ with $\mathbf{x}_1$) and note by optimality of the segmentation,

$$p_\epsilon(\mathcal{X}_i)\overline{\mathrm{F}}_\epsilon(p_\epsilon(\mathcal{X}_i) - \mu(\mathbf{x}_3)) > p_\epsilon(\mathcal{X}_j)\overline{\mathrm{F}}_\epsilon(p_\epsilon(\mathcal{X}_j) - \mu(\mathbf{x}_3)),$$
$$p_\epsilon(\mathcal{X}_j)\overline{\mathrm{F}}_\epsilon(p_\epsilon(\mathcal{X}_j) - \mu(\mathbf{x}_2)) > p_\epsilon(\mathcal{X}_i)\overline{\mathrm{F}}_\epsilon(p_\epsilon(\mathcal{X}_i) - \mu(\mathbf{x}_2)).$$

Combining these two inequalities gives

$$\frac{\overline{\mathrm{F}}_\epsilon(p_\epsilon(\mathcal{X}_i) - \mu(\mathbf{x}_3))}{\overline{\mathrm{F}}_\epsilon(p_\epsilon(\mathcal{X}_j) - \mu(\mathbf{x}_3))} > \frac{p_\epsilon(\mathcal{X}_j)}{p_\epsilon(\mathcal{X}_i)} > \frac{\overline{\mathrm{F}}_\epsilon(p_\epsilon(\mathcal{X}_i) - \mu(\mathbf{x}_2))}{\overline{\mathrm{F}}_\epsilon(p_\epsilon(\mathcal{X}_j) - \mu(\mathbf{x}_2))}.$$

Which can be further rearranged to $\frac{\overline{\mathrm{F}}_\epsilon(p_\epsilon(\mathcal{X}_i)-\mu(\mathbf{x}_3))}{\overline{\mathrm{F}}_\epsilon(p_\epsilon(\mathcal{X}_i)-\mu(\mathbf{x}_2))} > \frac{\overline{\mathrm{F}}_\epsilon(p_\epsilon(\mathcal{X}_j)-\mu(\mathbf{x}_3))}{\overline{\mathrm{F}}_\epsilon(p_\epsilon(\mathcal{X}_j)-\mu(\mathbf{x}_2))}$ which contradicts the PF2 property. Thus the optimal segmentation must be interval and $\mathcal{X}_i = \{\mathbf{x}|\mu(\mathbf{x}) \in [s_{i-1}, s_i)\}$ for some real numbers $s_i < s_{i+1}$.

(c) To show $p_\epsilon(\mathcal{X}_i) = p_\epsilon(\mathbf{x})$ for some $\mathbf{x} \in \mathcal{X}_i$, let $\mathbf{x}' = \arg\min_{\mathbf{x}\in\mathcal{X}_i} \mu(\mathbf{x})$ and recall $p_\epsilon(\mathcal{X}_i) = \arg\max \int_{\mu(\mathbf{x})\in[s_{i-1},s_i)} p\overline{\mathrm{F}}_\epsilon(p-s)f(\mu^{-1}(s))ds$. Now suppose $p_\epsilon(\mathcal{X}_i) < \mu(\mathbf{x}')$. By log-concavity, each function $\mathcal{R}_\epsilon(\mu(\mathbf{x}), p) := p\overline{\mathrm{F}}_\epsilon(p - \mu(\mathbf{x}))$ is unimodal, and thus increasing in p for $p \leq p_\epsilon(\mu(\mathbf{x}))$,

$$p_\epsilon(\mathcal{X}_i)\overline{\mathrm{F}}_\epsilon(p_\epsilon(\mathcal{X}_i) - s) \leq p_\epsilon(\mu(\mathbf{x}'))\overline{\mathrm{F}}_\epsilon(p_\epsilon(\mu(\mathbf{x}')) - s),$$

for any $s \in [s_i, s_{i+1}]$, which implies

$$\int_{\mu(\mathbf{x})\in[s_{i-1},s_i)} p_\epsilon(\mathcal{X}_i)\overline{\mathrm{F}}_\epsilon(p_\epsilon(\mathcal{X}_i) - s)f(\mu^{-1}(s))ds \leq \int_{\mu(\mathbf{x})\in[s_{i-1},s_i)} p_\epsilon(\mu(\mathbf{x}'))\overline{\mathrm{F}}_\epsilon(p_\epsilon(\mu(\mathbf{x}')) - s)f(\mu^{-1}(s))ds,$$

thus $p_\epsilon(\mathcal{X}_i) \geq p_\epsilon(\mathbf{x}')$. A symmetric argument similarly shows $p_\epsilon(\mathcal{X}_i) \leq \arg\max_{\mathbf{x}\in\mathcal{X}_i} p_\epsilon(\mathbf{x})$ for any $\mathcal{X}_i$. $\square$

Proof of Theorem 2. Suppose the firms prediction model $\mu(\mathbb{X})$ is supported on n values $\{x_i\}_{i=1}^n$, occurring with probabilities $\{q_i\}_{i=1}^n$, where $x_1 \leq x_2 \leq \ldots \leq x_n$. By Lemma 2, the optimal segmentation can be indexed by the sequence $\{s_i\}_{i=0}^k$ which is contained in the support of $\mu(\mathbb{X})$. Let the optimal price for segment $[s_{i-1}, s_i)$ be $p_\epsilon([s_{i-1}, s_i))$ i.e.

$$p_\epsilon([s_{i-1}, s_i)) = \arg\max_{p_i} p_i \Pr(\mu(\mathbf{x}) + \epsilon \geq p_i | \mu(\mathbf{x}) \in [s_{i-1}, s_i)) \Pr(\mu(\mathbf{x}) \in [s_{i-1}, s_i)).$$

We wish to find $\{s_i\}_{i=0}^k \subset \{x_i\}_{i=1}^n$ that maximizes

$$\sum_{i=1}^k p_\epsilon([s_{i-1}, s_i)) \Pr(\mu(\mathbf{x}) + \epsilon \geq p_\epsilon([s_{i-1}, s_i)) | \mu(\mathbf{x}) \in [s_{i-1}, s_i)) \Pr(\mu(\mathbf{x}) \in [s_{i-1}, s_i)).$$

We suppose the time to compute $p_\epsilon(s_i, s_{i+1})$ for any segment $[s_i, s_{i+1})$ is upper bounded by m_ϵ. Now note there are at most $\frac{n(n+1)}{2}$ intervals to consider, and we can create a table to store the optimal prices for all possible intervals in $\mathrm{O}(n^2 m_\epsilon)$ time.

We now give a dynamic programming solution that uses time $O(kn^2)$ and to populate a table of size kn. Define $D[n', k']$ as the optimal k'-market segmentation that considers only the n' lowest predicted valuations $\{(x_i, q_i)\}_{i=1}^{n'}$, our goal is to compute $D[n, k]$ which is the revenue of the optimal FBMSP (the optimal policy can further be reconstructed by standard backward search). Our algorithm depends on the following observation: consider the optimal k-market segmentation and suppose $[s_{k-1}, s_k] = [x_{i_k}, x_n]$ defines the k^{th} segment. If one considers the market without the customers in the k^{th} segment, the remaining $k-1$ segments must be an optimal $(k-1)$-market segmentation on $\{(x_i, q_i)\}_{i=1}^{i_k - 1}$. Formally, we express this observation as the following recursion,

$$D[n', k'] = \max_{l \in [n'-1]} D[l, k'-1] + p_\epsilon([s_{i-1}, s_i)) \sum_{i=l+1}^{n'} \Pr(\mu(\mathbf{x}) + \epsilon \geq p_\epsilon([s_{i-1}, s_i)) | \mu(\mathbf{x}) \in [s_{i-1}, s_i)) q_i, \tag{23}$$

which states that the optimal k'-market segmentation on the lowest n' valuations, is equal to some optimal $(k'-1)$-segmentation on a smaller market, plus the value of the k'^{th} segment. Using Eq. (23) we may populate a table of size kn, starting at $D[0, 0] = 0$, and computing column-wise. The maximization in Eq. (23) takes at most $n' - 1$ calculations. If the optimal price and revenue for each segment are stored before the iteration, the dynamic programming can be finished in $\mathrm{O}(kn^2)$ time. Thus, the optimal feature-based market segmentation can be computed in $\mathrm{O}(n^2(k + m_\epsilon))$ time.

$\square$

A.2.3 Omitted Proofs from Section 2.4

Proof of Theorem 3. Let $\{s_i\}_{i=0}^k \in \mathbb{R}^{k+1}$ denote an optimal interval k-market segmentation for $\mathcal{R}_{kXP}^{\mu(\mathbb{X})}$. Consider the sub-optimal feature-based market segmentation which uses segments $\mathcal{X}_i = \{\mathbf{x} \,|\, \mu(\mathbf{x}) \in [s_{i-1}, s_i)\}$ and prices p_i. Note for all $\mathbf{x} \in \mathcal{X}_i$, $p_i \Pr(\mu(\mathbf{x}) + \epsilon \geq p_i) \geq \mathcal{R}_\epsilon(s_{i-1})$ and thus summing over all segments

$$\mathcal{R}_{kXP}^{\mu(\mathbb{X})+\epsilon} \geq \sum_{i=1}^k \int_{\mu(\mathbf{x}) \in [s_{i-1}, s_i)} p_i \Pr(\mu(\mathbf{x}) + \epsilon \geq p_i)\, f_{\mathbb{X}}(\mathbf{x}) dx \geq \sum_{i=1}^k \mathcal{R}_\epsilon(s_{i-1}) \int_{\mu(\mathbf{x}) \in [s_{i-1}, s_i)} f_{\mathbb{X}}(\mathbf{x}) d\mathbf{x} \tag{24}$$

Now,

$$\begin{aligned}
\mathbb{E}_{\mathbb{X}\sim F_{\mathbb{X}}}[\mathcal{R}_\epsilon(\mu(\mathbb{X}))] - \mathcal{R}_{kXP}^{\mu(\mathbb{X})+\epsilon} &\leq \sum_{i=1}^k \int_{\mu(\mathbf{x}) \in [s_{i-1}, s_i)} (\mathcal{R}_\epsilon(\mu(\mathbf{x})) - \mathcal{R}_\epsilon(s_{i-1}))\, f_{\mathbb{X}}(\mathbf{x}) d\mathbf{x} && \text{Eq. (24)} \\
&\leq \sum_{i=1}^k \int_{\mu(\mathbf{x}) \in [s_{i-1}, s_i)} (\mu(\mathbf{x}) - s_{i-1}) \overline{\mathrm{F}}_\epsilon(\theta_\epsilon(\mu(\mathbf{x}))) f_{\mathbb{X}}(\mathbf{x}) d\mathbf{x} && \text{Lemma 1(b)} \\
&\leq \sum_{i=1}^k \int_{\mu(\mathbf{x}) \in [s_{i-1}, s_i)} (\mu(\mathbf{x}) - s_{i-1}) f_{\mathbb{X}}(\mathbf{x}) d\mathbf{x} \\
&= \mathbb{E}_{X\sim F_{\mathbb{X}}}[\mu(\mathcal{X}_i)] - \mathcal{R}_{kXP}^{\mu(\mathcal{X}_i)}
\end{aligned}$$

as desired. $\square$

Proof of Theorem 4. Let $L = \inf_{\mathbf{x}} \mu(\mathbf{x})$ and $U = \sup_{\mathbf{x}} \mu(\mathbf{x})$, and recalling the proof of Theorem 3 we have,

$$\mathbb{E}_{V\sim F}[\mathcal{R}_\epsilon(V)] - \mathcal{R}_{kXP}^V \leq \sum_{i=1}^k \int_{\mu(\mathbf{x}) \in [s_{i-1}, s_i)} (\mathcal{R}_\epsilon(\mu(\mathbf{x})) - \mathcal{R}_\epsilon(s_{i-1}))\, \overline{\mathrm{F}}_\epsilon(\theta_\epsilon(s_{i-1}))\, f_{\mathbb{X}}(\mathbf{x}) d\mathbf{x}.$$

Now instead we choose the sub-optimal interval segmentation such that it subdivides the quantile space of $F_{\mathbb{X}}$ into k equal regions i.e., $\{s_i\}_{i=0}^{k}$ such that $\overline{\mathrm{F}}(s_{i-1}) - \overline{\mathrm{F}}(s_i) = \frac{1}{k}$ for $i = 1, ..., k$. Then

$$
\begin{aligned}
\mathbb{E}_{\mu(\mathbb{X})\sim F}[\mathcal{R}_{\epsilon}(\mu(\mathbb{X}))] - \mathcal{R}_{kXP}^{\mu(\mathbb{X})+\epsilon} &\leq \sum_{i=1}^{k} \int_{\mu(\mathbf{x})\in[s_{i-1},s_i)} \left(\mathcal{R}_{\epsilon}(\mu(\mathbf{x})) - \mathcal{R}_{\epsilon}(s_{i-1})\right) f_{\mathbb{X}}(\mathbf{x}) d\mathbf{x} \\
&\leq \sum_{i=1}^{k} \int_{\mu(\mathbf{x})\in[s_{i-1},s_i)} (\mu(\mathbf{x}) - s_{i-1})\overline{\mathrm{F}}_{\epsilon}(\theta_{\epsilon}(\mu(\mathbf{x})))f_{\mathbb{X}}(\mathbf{x})d\mathbf{x} \quad \text{Lemma 1(b)} \\
&\leq \sum_{i=1}^{k} \int_{\mu(\mathbf{x})\in[s_{i-1},s_i)} (\mu(\mathbf{x}) - s_{i-1})f_{\mathbb{X}}(\mathbf{x})d\mathbf{x} \\
&= \sum_{i=0}^{k} (s_i - s_{i-1}) \int_{\mu(\mathbf{x})\in[s_{i-1},s_i)} f_{\mathbb{X}}(\mathbf{x})d\mathbf{x} \\
&= \sum_{i=1}^{k} (s_i - s_{i-1}) \left(\overline{\mathrm{F}}(s_{i-1}) - \overline{\mathrm{F}}(s_i)\right) \\
&= \frac{\sum_{i=1}^{k} (s_i - s_{i-1})}{k} \\
&= \frac{U - L}{k},
\end{aligned}
$$

where the third inequality follows $\overline{\mathrm{F}}_{\epsilon}\left(\theta_{\epsilon}(\cdot)\right) \leq 1$, the third equality comes from the choice of $\{s_i\}_{i=0}^{k}$. Finally, summing $s_i - s_{i-1}$ we get the final equality as desired. □

Proof of Theorem 5. Fix some $k \geq 2$, we will prove the rearranged inequality $\mathcal{R}_{(k-1)XP} + \mathcal{R}_{(k+1)XP} \leq 2\mathcal{R}_{kXP}$ by explicitly constructing feasible (but not necessarily optimal) size k segmentations. Note, since ϵ is log-concave, by Lemma 2 the optimal segmentation for any k is interval, and can be described by the sequence of numbers $\{s_i^k\}_{i=0}^{k}$ such that $\mathcal{X}_i^k = \{\mathbf{x}|\mu(\mathbf{x}) \in [s_{i-1}, s_i)\}$. Further, let $\mathcal{S}_{k-1} := \{s_i^{k-1}\}_{i=0}^{k-1}$, $\mathcal{S}_k := \{s_i^k\}_{i=0}^{k}$ and $\mathcal{S}_{k+1} := \{s_i^{k+1}\}_{i=0}^{k+1}$ be the optimal segmentations of size $k-1$, k, and $k+1$, respectively, as described by segmentation endpoints. Note by definition $s_{k-1}^{k-1} = s_k^k = s_{k+1}^{k+1} = \sup_{\mathbf{x}} \mu(\mathbf{x})$. Our proof will proceed in two cases.

In this case, the first segment of the optimal $(k+1)$ segmentation is before the first segment of the $(k-1)$ segmentation, see Fig. 17 for an illustration. Consider two feasible k segmentation $\mathcal{S}_k' = s_0^{k+1} \cup \mathcal{S}_{k-1}$ and $\mathcal{S}_k'' = \mathcal{S}_{k+1} \setminus s_0^{k+1}$, see Fig. 18 for another illustration.

Now note the combined revenue from $\mathcal{S}'_k$ and $\mathcal{S}''_k$ fully covers the revenue from $\mathcal{S}_{k-1}$ (as a subset of $\mathcal{S}'_k$) and $\mathcal{S}_{k+1}$ except the revenue from the first segment (all other segments of the optimal $(k+1)$ segmentation are covered by $\mathcal{S}''_k$). Now, in the constructed segmentation $\mathcal{S}'_k$, the unaccounted for first segment has end points $[s_0^{k+1}, s_0^{k-1})$, which by assumption contains the first segment of the $(k+1)$ segmentation $[s_0^{k+1}, s_1^{k+1})$. Note if one segment subsumes another, it provides more revenue, i.e., if $\mathcal{X}_i \subset \mathcal{X}_j$ then

$$\max_p p \Pr\left(\mu\left(\mathbf{x}\right) + \epsilon \geq p | \mathbf{x} \in \mathcal{X}_i\right) \Pr(\mathcal{X}_i) \leq \max_p p \Pr\left(\mu\left(\mathbf{x}\right) + \epsilon \geq p | \mathbf{x} \in \mathcal{X}_j\right) \Pr(\mathcal{X}_j). \tag{25}$$

and thus Eq. (25) implies $\mathcal{R}_{(k-1)XP} + \mathcal{R}_{(k+1)XP} \leq \mathcal{R}_{kXP}(\mathcal{S}'_k) + \mathcal{R}_{kXP}(\mathcal{S}''_k) \leq 2\mathcal{R}_{kXP}$.

Figure 17: Case 1 of Theorem 5.

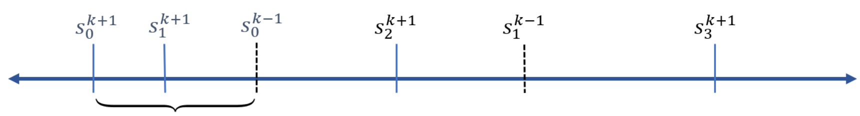

Note. Depicted are the first segment end points of an interval segmentation of size k+1 (solid line) and k-1 (dashed line). The first segment $\left[s_0^{k+1}, s_1^{k+1}\right)$ in the $(k+1)$ segmentation is before the first segment $\left[s_0^{k-1}, s_1^{k-1}\right)$ of the $(k-1)$ segmentation.

Figure 18: New feasible k-segmentations for case 1 of Theorem 5.

Note. Depicted are the constructed size k interval segmentations. In the left panel $\mathcal{S}'_k$ is shown which is equal to the optimal $(k-1)$ segmentation plus a new first segment $\left[s_0^{k+1}, s_0^{k-1}\right)$. In the right panel $\mathcal{S}''_k$ is shown which is equal to the optimal $(k+1)$ segmentation with the first segment $\left[s_0^{k+1}, s_1^{k+1}\right)$ removed.

In this case, there is an i such that segment $i+1$ of the optimal $(k+1)$ segmentation that is subsumed by segment i of the optimal $(k-1)$ segmentation. As in Case 1, we will construct feasible k segmentations assuming the condition of Case 2 holds. Before constructing the feasible segmentations, we will need two simple facts, both of which follow from Lemma 2.

$$\textbf{Fact 1: } \text{If } \mathcal{X}_1 = [s_1, s_2), \mathcal{X}_2 = [s_1, s_2 + \Delta), \text{where } \Delta \geq 0, \text{then } p_\epsilon(\mathcal{X}_2) \geq p_\epsilon(\mathcal{X}_1), \tag{26}$$

$$\textbf{Fact 2: } \text{If } \mathcal{X}_1 = [s_1, s_2), \mathcal{X}_2 = [s_1 - \Delta, s_2), \text{where } \Delta \geq 0, \text{then } p_\epsilon(\mathcal{X}_2) \leq p_\epsilon(\mathcal{X}_1). \tag{27}$$

Fix the i such that $s_{i-1}^{k-1} \leq s_i^{k+1} \leq s_{i+1}^{k+1} \leq s_i^{k-1}$. Such an arrangement of segmentation points is shown in Fig. 19. Now define the feasible k-segmentations as

$$\mathcal{S}_k' = \left\{s_j^{k-1}\right\}_{j=1}^{i-1} \cup \left\{s_j^{k+1}\right\}_{j=i+1}^{k+1},$$

and

$$\mathcal{S}_k'' = \left\{s_j^{k+1}\right\}_{j=1}^{i} \cup \left\{s_j^{k-1}\right\}_{j=i}^{k-1}.$$

Each new arrangement of segmentation points consists of splicing the beginning of the $(k-1)$ segmentation with the end of the $(k+1)$ segmentation, or vice versa, with the middle segment added or removed. An example of such a segmentation construction is shown in Fig. 20. Note compared to the $(k-1)$ and $(k+1)$ segmentations, the new segment in $\mathcal{S}_k'$ is $\left[s_{i-1}^{k-1}, s_{i+1}^{k+1}\right)$, and the new segment in $\mathcal{S}_k''$ is $\left[s_i^{k+1}, s_i^{k-1}\right)$, and all the other segments are the same as segments in the optimal $k-1$ or $(k+1)$ segmentation. The only segments in the optimal $(k-1)$ or $(k+1)$ segmentation unaccounted for (i.e. not contained in $\mathcal{S}_k'$ or $\mathcal{S}_k''$) are $[s_{i-1}^{k-1}, s_i^{k-1})$ and $[s_i^{k+1}, s_{i+1}^{k+1})$. Now let $p_1 := p_\epsilon\left([s_{i-1}^{k-1}, s_i^{k-1})\right)$, $p_2 := p_\epsilon\left([s_i^{k+1}, s_{i+1}^{k+1})\right)$ be the optimal prices on the those unaccounted for segments, we will need to determine the prices for each new segments in $\mathcal{S}_k'$ and $\mathcal{S}_k''$ such that the combined revenue from them fully covers the revenue from the unaccounted for segments.

Note that it is unclear which of p_1 and p_2 is larger, we will argue in two sub-cases based on their ordering. Suppose we have $p_1 \leq p_2$, and let the price for new segments in $\mathcal{S}_k'$ and $\mathcal{S}_k''$ be

$$p_\epsilon\left(\left[s_{i-1}^{k-1}, s_{i+1}^{k+1}\right)\right) = p_1, \quad p_\epsilon\left(\left[s_i^{k+1}, s_i^{k-1}\right)\right) = p_2.$$

Let $\mathcal{R}_\epsilon([s_{j-1}, s_j), p)$ be the revenue from segment $[s_{j-1}, s_j)$ when the price on that segment is p. The difference in revenue between the new and unaccounted for segments is then,

$$
\begin{aligned}
&\mathcal{R}_\epsilon\left(\left[s_{i-1}^{k-1}, s_{i+1}^{k+1}\right), p_1\right) + \mathcal{R}_\epsilon\left(\left[s_{i}^{k+1}, s_{i}^{k-1}\right), p_2\right) - \mathcal{R}_\epsilon\left(\left[s_{i-1}^{k-1}, s_{i}^{k-1}\right), p_1\right) - \mathcal{R}_\epsilon\left(\left[s_{i}^{k+1}, s_{i+1}^{k+1}\right), p_2\right) \\
&= \left(\mathcal{R}_\epsilon\left(\left[s_{i-1}^{k-1}, s_{i+1}^{k+1}\right), p_1\right) - \mathcal{R}_\epsilon\left(\left[s_{i-1}^{k-1}, s_{i}^{k-1}\right), p_1\right)\right) + \left(\mathcal{R}_\epsilon\left(\left[s_{i}^{k+1}, s_{i}^{k-1}\right), p_2\right) - \mathcal{R}_\epsilon\left(\left[s_{i}^{k+1}, s_{i+1}^{k+1}\right), p_2\right)\right) \\
&= \mathcal{R}_\epsilon([s_{i+1}^{k+1}, s_i^{k-1}), p_2) - \mathcal{R}_\epsilon([s_{i+1}^{k+1}, s_i^{k-1}), p_1),
\end{aligned}
$$

where the second equality follows from the fact that $\left[s_{i-1}^{k-1}, s_{i+1}^{k+1}\right) \subset \left[s_{i-1}^{k-1}, s_{i}^{k-1}\right)$, and $\left[s_{i}^{k+1}, s_{i+1}^{k+1}\right) \subset \left[s_{i}^{k+1}, s_{i}^{k-1}\right)$. Since every element in $\left[s_{i}^{k+1}, s_{i+1}^{k+1}\right)$ is less than any element in $\left[s_{i+1}^{k+1}, s_{i}^{k-1}\right)$, by Eq. (26) and Eq. (27), we have the price dominance

$$p_1 \leq p_2 \leq p_\epsilon([s_{i+1}^{k+1}, s_i^{k-1})),$$

which implies

$$\mathcal{R}_\epsilon([s_{i+1}^{k+1}, s_i^{k-1}), p_2) \geq \mathcal{R}_\epsilon([s_{i+1}^{k+1}, s_i^{k-1}), p_1).$$

Thus, $\mathcal{R}_{(k-1)XP} + \mathcal{R}_{(k+1)XP} \leq \mathcal{R}_{kXP}(\mathcal{S}'_k) + \mathcal{R}_{kXP}(\mathcal{S}''_k) \leq 2\mathcal{R}_{kXP}$. In the second sub-case when $p_2 \leq p_1$, the proof follows symmetrically now using p_2 for the price of the new segment in $\mathcal{S}'_k$ and p_1 for the price in the new segment for $\mathcal{S}''_k$, we omit it for brevity.

Figure 19: Case 2 of Theorem 5.

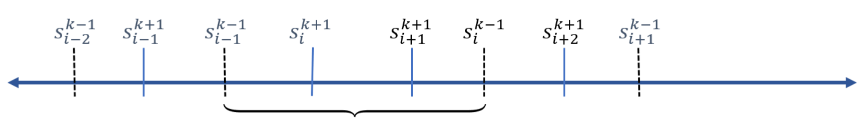

Note. Depicted are the $i+1$ segment end points of the an interval segmentation of size k+1 (solid line) and i segment of size k-1 segmentation (dashed line). The $i+1^{th}$ segment $\left[s_{i}^{k+1}, s_{i+1}^{k+1}\right)$ of $(k+1)$ segmentation is fully contained in the i^{th} segment $\left[s_{i-1}^{k-1}, s_{i}^{k-1}\right)$ of $(k-1)$ segmentation.

To complete the proof, we now must show that Case 1 and Case 2 are the only cases i.e., if the condition for Case 1 does not hold, the condition for Case 2 must hold for some i.

Figure 20: New feasible k-segmentations for case 2 of Theorem 5.

Note. Depicted are the constructed size k interval segmentations. The new k-segmentations are constructed by crossing over $(k-1)$ segmentation and $(k+1)$ segmentation at s_{i-1}^{k-1}. In the left panel, $\mathcal{S}'_k$ is shown; before s_{i-1}^{k-1}, it contains the segments of $\mathcal{S}_{k-1}$, after s_{i+1}^{k+1}, it contains segments of $\mathcal{S}_{k+1}$. In the right panel, $\mathcal{S}''_k$ is shown; before s_i^{k+1}, it contains the segments of $\mathcal{S}_{k+1}$, after s_i^{k-1}, it contains the segments of $\mathcal{S}_{k-1}$.

To see this, imagine Case 1 does not hold, then $s_1^{k+1} > s_0^{k-1}$. If then $s_2^{k+1} < s_1^{k-1}$ the proof is complete, so assume $s_2^{k+1} \geq s_1^{k-1}$. If then $s_3^{k+1} < s_2^{k-1}$ the proof is complete and so on. Iterating, since $s_{k+1}^{k+1} = s_{k-1}^{k-1}$ the sequence of deductions must terminate at some i for which Case 2 holds. Thus Case 1 and Case 2 cover all cases, which completes the proof.

$\square$

A.3 Transforming a Probit Regression Model into a Valuation Model

In this section we overview how to transform a prediction model for the sales probability into a linear valuation model. In practice we cannot observe a customer's valuation for one product directly. Instead, we observe whether the customer will buy the product or not, at the offered price p (see [30] for more details). Assume that the unobserved continuous dependent variable Y is the customer's true valuation or willingness to pay (WTP) for the product. Further, suppose the relation of Y and customer's feature X is

$$Y = \beta_0 + X\beta + \epsilon, \tag{28}$$

where $\epsilon \sim N(0, \sigma)$. Customer i's decision I_i will be

$$I_i = \begin{cases} 1, & \text{if } Y_i \geq p_i, \\ 0, & \text{otherwise,} \end{cases} \tag{29}$$

where p_i is the price offered to customer i. Then, the probability that customer i with features X_i will buy the product is

$$\begin{aligned} \Pr(I_i = 1) = \Pr(Y_i \geq p) &= \Pr(\beta_0 + X_i\beta + \epsilon_i \geq p) \\ &= \Pr\left(\frac{\epsilon_i}{\sigma} \geq \frac{p - X_i\beta}{\sigma}\right) \\ &= 1 - \Phi\left(\frac{p - \beta_0 - X_i\beta}{\sigma}\right), \end{aligned}$$

where Φ is the cumulative distribution function for the standard normal distribution. Using p and X as explanatory variables, the probit regression model is then

$$\Pr(I = 1|X) = 1 - \Phi(\beta_0' + p\beta_p + X\beta').$$

Therefore, we can use the maximum likelihood estimator (MLE) of the probit regression model to recover the regression model of customer's valuation, i.e.,

$$\hat{\beta}_0 = \frac{\hat{\beta}_0'}{\hat{\beta}_p}, \quad \hat{\beta} = \frac{\hat{\beta}'}{\hat{\beta}_p}, \quad \hat{\sigma} = \frac{1}{\hat{\beta}_p},$$

where $\hat{\beta}_0'$, $\hat{\beta}_p$, $\hat{\beta}'$ are the MLE of β_0', β_p, β'. Further, the regression model of customer's valuation recovered from probit regression model is asymptotically unbiased if the price variation is large enough.

A.4 Constant Factor Approximation for General Error

In this section we will describe how to obtain a 1-1/e approximation of the optimal FBMSP when the residuals are independent and follow an arbitrary distribution, as sketched in Remark 1. Through the section we will use 2^V to represent the power set of some set V.

$(1-1/e)$ **Approximation Algorithm**: Our polynomial time approximation algorithm will follow from the *submodularity* of the objective function for FBMSP, defined as follows:

Definition 7 (Submodularity)**.** *A set function $f : 2^V \to \mathbb{R}$ is submodular if for every $A, B \subseteq V$,*

$$f(A \cap B) + f(A \cup B) \leq f(A) + f(B).$$

An important subclass of submodular functions are those which are monotone, i.e., functions for which enlarging the choice set cannot cause the function value to decrease.

Definition 8 (Monotonicity)**.** *A set function $f : 2^V \to \mathbb{R}$ is monotone if for every $A \subseteq B \subseteq V$, $f(A) \leq f(B)$.*

We will show that the objective function for FBMSP can be expressed as a set function over prices which is monotone and submodular. Note that for n customers with predicted valuations $\{\mu(x_i)\}_{i=1}^n$ and for error distribution ϵ supported on m points, there are at most $\mathrm{O}(nm)$ distinct possible valuation realizations. Further, any optimal price for a segment must correspond to one of these realizations (since if not, raising the price until it reaches a valuation in the support is strictly revenue improving). Thus the set of potential prices is a polynomially sized set equivalent to the set of potential realized valuations, and the revenue objective of FBMSP can be viewed as a set function over a size k subset of that price set.

Specifically, if f is the revenue function of FBMSP on price set A, it takes the form,

$$f(A) = \sum_{i=1}^{n} \max_{p \in A} p\overline{\mathrm{F}}(p - \mu(x_i)).$$

Then expressed as a set function over the prices, optimal FBMSP is the solution to

$$\max_{|A| \leq k} \sum_{i=1}^{n} \max_{p \in A} p\overline{\mathrm{F}}(p - \mu(x_i)).$$

The monotonicity of the revenue objective is easy to see since, by definition, enlarging the set of possible prices that can be used for a segment will keep at least the same revenue as for a smaller set of prices. The submodularity comes from the fact that any customer i facing the prices in price set $A \cap B$ will result in less revenue than when facing the prices in price set A or B, i.e.,

$$\max_{p \in A \cap B} p\overline{\mathrm{F}}(p - \mu(x_i)) \leq \min \left\{ \max_{p \in A} p\overline{\mathrm{F}}(p - \mu(x_i)), \quad \max_{p \in B} p\overline{\mathrm{F}}(p - \mu(x_i)) \right\},$$

further,

$$\max_{p \in A \cup B} p\overline{\mathrm{F}}(p - \mu(x_i)) = \max \left\{ \max_{p \in A} p\overline{\mathrm{F}}(p - \mu(x_i)), \quad \max_{p \in B} p\overline{\mathrm{F}}(p - \mu(x_i)) \right\}.$$

Note that

$$\begin{aligned} \max_{p \in A} p\overline{\mathrm{F}}(p - \mu(x_i)) + \max_{p \in B} p\overline{\mathrm{F}}(p - \mu(x_i)) = \min \left\{ \max_{p \in A} p\overline{\mathrm{F}}(p - \mu(x_i)), \quad \max_{p \in B} p\overline{\mathrm{F}}(p - \mu(x_i)) \right\} + \\ \max \left\{ \max_{p \in A} p\overline{\mathrm{F}}(p - \mu(x_i)), \quad \max_{p \in B} p\overline{\mathrm{F}}(p - \mu(x_i)) \right\}, \end{aligned}$$

and thus combining these observations and summing over all customers proves submodularity of the objective function.

Note that positive monotone submodular function maximization with cardinality constraints is NP-hard in general (see [84]). The cardinality constraint in FBMSP is the number of segments (the same as the number of prices). [100] shows that a greedy algorithm can obtain an approximation guarantee of $(1 - 1/e)$ for class of monotone submodular functions with cardinality constraints. Since FBMSP problem is can be written as a problem of maximizing a monotone submodular function with cardinality constraints, it can be approximated to at least within a factor of $(1 - 1/e)$ via the same greedy algorithm.

A.5 Omitted Figures

Figure 21: Prediction of the probit regression model.

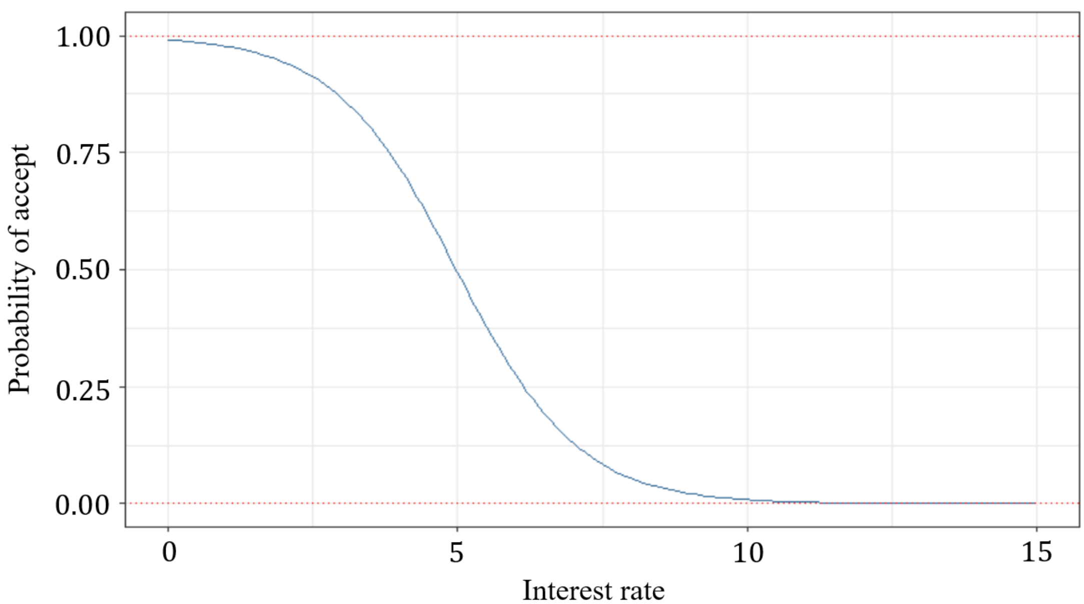

Note. Depicted is the output of a probit regression model to predict the probability of mortgage acceptance, our proxy for purchase in the loan setting. The model is trained using features in Table 2.

Figure 22: An example of the error distribution $\overline{\mathrm{F}}_\epsilon$, constructed for the proof of Theorem 1.

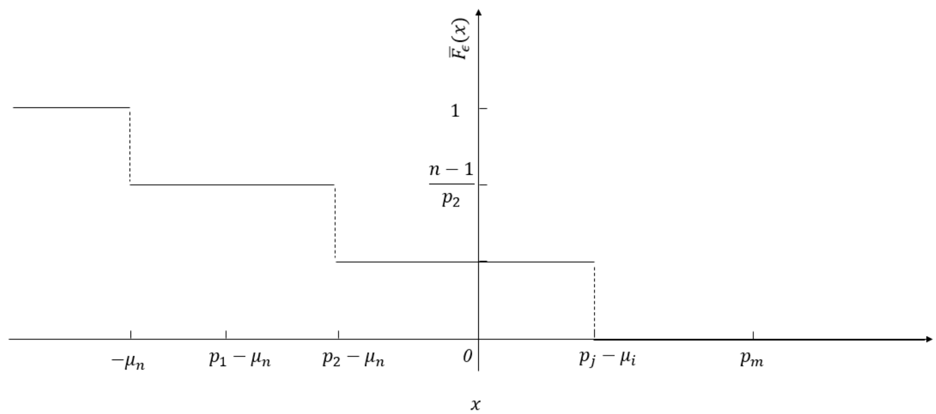

Note. Depicted is an example of the cCDF $\overline{\mathrm{F}}_\epsilon$ constructed to prove the hardness of FBMSP. Note on the x-axis are the valuation support points, and that the resultant error distribution is a step-function on these supports.

Appendix B Pricing for Online Dating

B.1 Omitted Examples

Example 4 (Properties of Period L Pricing)**.** *In this example, we demonstrate three properties of period L pricing. First, we show that the optimal price for period L pricing may not be unique, with potentially many locally/globally optimal prices, and thus we cannot rely on first order conditions alone in our analysis for general period L pricing. Second, we show that the profit of period L pricing may not be monotone over L, even for MHR valuation distributions. Third, we show that for any given period length L, there exists model parameters such that L is the optimal period length precluding the existence of a universally optimal choice of L.*

To demonstrate the non-uniqueness of the optimal prices for period L pricing, suppose a user's valuations are uniformly distributed on the range $[0, 1]$*, i.e.,* $V \sim U[0, 1]$*, and suppose the cost* $c = 0.1$*, the match rate* $q = 0.8$*, the valuation decays as* $\delta = 0.25$*, the total time users can stay on the platform* $T = 1$*, and the period length* $L = T/2 = 0.5$*. Then referring to Fig. 23 left panel, we can see there are two locally optimal subscription prices, one is near* 0.12*, the other near* 0.15*. To understand this phenomenon, note in this example that users can stay on the platform for at most two periods. Therefore, the online dating platform can set the price to be high, which will force users to stay on the platform for at most one period, or the online dating platform can set the price to be low, which will allow users stay for up to two periods. Each instance corresponds to a different locally optimal price.*

To demonstrate the non-monotonicity of $\mathcal{R}(L)$ *as a function of* L*, consider a market similar to the one described above where users' valuations are distributed as* $U[0, 1]$*,* $c = 0.05$*,* $q = 0.2$*, and now valuations decay as* $\delta = 0.8$*, and* $T = 50$*. In such a market, the optimal choice of L can be numerically computed to be 4.55, and achieves profit 0.40. In the middle panel of Fig. 23, we see the non-monotonicity of profits in L, and further note that there also exists locally (but not globally) optimal choices of period L before the optimal period length L.*

Finally, to demonstrate that every choice of L can be optimal given the right market parameters, again consider a market where user's valuations $V \sim U[0,1]$, $q = 0.2$, $\delta = 0.8$, $T = 50$. In the right panel of Fig. 23 we vary the cost c and compute the optimal choice of L via grid search. We note that L is decreasing in a smooth fashion as c increases.

Figure 23: Examples demonstrating properties of LP.

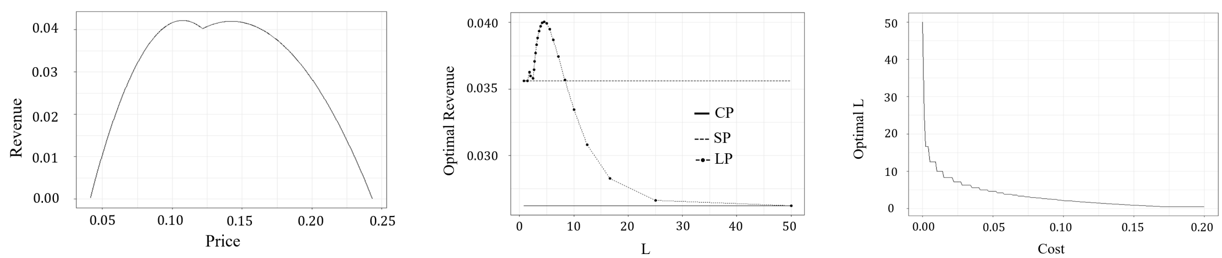

Note. In the left panel, depicted are an online dating platform's profit vs. subscription prices, where user's valuations $V \sim U[0,1]$, unit operating cost is $c = 0.1$, match rate is $q = 0.8$, user's patience factor is $\delta = 0.25$, total market size is $T = 1$, and the period length is $L = T/2 = 0.5$. In the middle panel, depicted are an online dating platform's profit using LP, SP and CP, as period length L varies, where user's valuations $V \sim U[0,1]$, unit operating cost $c = 0.05$, match rate $q = 0.2$, user's patience factor $\delta = 0.8$, and total time $T = 50$. In the right panel, depicted is the optimal period length L when cost c changes, where user's valuations are $V \sim U[0,1]$, the match rate is $q = 0.2$, the user's patience factor is $\delta = 0.8$, and the market size is $T = 50$.

Example 5 (Tightness of Lemma 6)**.** *In this example, we give an instance of our model such that Lemma 6 is tight. Specifically, suppose valuations are fixed and drawn from a point-mass distribution on v, let $c = 0$, $T = \infty$, and let q and δ be arbitrary. In this case the optimal short period price is $p^* = vq\left(\frac{\log(\delta)-q}{\log(\delta)}\right)^{\frac{\log(\delta)}{q}}$, and the profit of optimal short period pricing is,*

$$\mathcal{R}_{SP} = \left(\frac{vq}{q - \log(\delta)}\right)\left(\frac{\log(\delta) - q}{\log(\delta)}\right)^{\frac{\log(\delta)}{q}}.$$

Similarly, the profit of optimal long period pricing is,

$$\mathcal{R}_{CP} = \left(\frac{vq}{q - \log(\delta)}\right).$$

The ratio between $\mathcal{R}_{SP}$ and $\mathcal{R}_{CP}$ is then,

$$\frac{\mathcal{R}_{SP}}{\mathcal{R}_{CP}} = \left(1 - \frac{q}{\log(\delta)}\right)^{\frac{\log(\delta)}{q}},$$

therefore, the approximation ratio is tight.

Example 6 (Only SP can Achieve Universal Constant Factor Approximation)**.** *In this example, we show that no other fixed pricing strategy besides short period pricing can guarantee any constant factor approximation for all parameters. Suppose user's valuations follow a unit mass distribution of v , i.e., $V = v$ with probability 1, and $T = \infty$. For any fixed period L pricing, when $c = vq^2\left(1 - \delta^L e^{-qL}\right) / \left((q - \log(\delta))\left(1 - e^{-qL}\right)\right)$, by Eq. (7), to achieve positive profit the price of period L pricing must satisfy $p \geq c\left(1 - e^{-qL}\right)/q = vq\left(1 - \delta^L e^{-qL}\right)/(q - \log(\delta))$. Note in Eq. (4), a user will pay the price for first period if and only if*

$$\frac{vq\left(1 - \delta^L e^{-qL}\right)}{q - \log(\delta)} \geq p.$$

Thus period L pricing cannot earn any profit when $c = vq^2\left(1 - \delta^L e^{-qL}\right) / \left((q - \log(\delta))\left(1 - e^{-qL}\right)\right)$ whereas short period pricing can still earn some profit by setting $p \geq c$ but $p < vq$, which is always possible since $\frac{1-\delta^L e^{-qL}}{1-e^{-qL}}$ is decreasing in L, and $\frac{q\left(1-\delta^L e^{-qL}\right)}{q-\log(\delta)} < 1$ for any $L > 0$.

Example 7 (Profit and Match Proportion when $\delta = 1$)**.** *In this example, we give an instance of our model such that the profit dominance in Theorem 7 is tight. Specifically, let $\delta = 1$ and let F, c, q, and T be arbitrary. In this case the profit of SP and CP can be written as,*

$$\mathcal{R}_{SP}(p) = \int_0^T (p-c)\overline{F}(pq^{-1})e^{-qt}dt = (p-c)\overline{F}(pq^{-1})\int_0^T e^{-qt}dt,$$

$$\mathcal{R}_{CP}(p) = \int_0^T \left(\frac{pq}{1-e^{-qT}} - c\right)\overline{F}\left(\frac{p}{1-e^{-qT}}\right)e^{-qt}dt = \left(\frac{pq}{1-e^{-qT}} - c\right)\overline{F}\left(\frac{p}{1-e^{-qT}}\right)\int_0^T e^{-qt}dt.$$

Let $p_S^* = \arg\max_p \mathcal{R}_{SP}(p)$, $p_C^* = \arg\max_p \mathcal{R}_{CP}(p)$, then, solving the optimization for SP and CP, we obtain the optimal prices $\left(1 - e^{-qT}\right)p_S^* = qp_C^*$, and moreover $\mathcal{R}_{SP} = \mathcal{R}_{CP}$.

Example 8 (Manipulation Gap)**.** *In this example, we consider SP and CP when the platform has access to two types of potential matches it can display. Suppose valuations are fixed and drawn from a point-mass distribution on* v*, let* $c = 0$*,* $T = 2$*,* $\delta = 1$*, and suppose* $q_1 = 0$*,* $q_2 = 1$*, and* $t_1 = t_2 = 1$*. Then the match rate perceived by users is* $q = \frac{q_1+q_2}{2} = 0.5$*. The profit of long period pricing that chooses to show type 2 potential matches first is,*

$$\mathcal{R}_{CP}(\{q_2, q_1\}) = v\left(1 - e^{-qT}\right) = v\left(1 - \frac{1}{e}\right),$$

The profit of short period pricing that chooses to show type 1 potential matches first is,

$$\mathcal{R}_{SP}(\{q_1, q_2\}) = vq\left(\frac{1 - e^{-q_1}}{q_1} + e^{-q_1}\left(\frac{1 - e^{-t_2 q_2}}{q_2}\right)\right) = v\left(1 - \frac{1}{2e}\right).$$

By Example 7, when the match rate was homogeneous the profit of SP and CP was the same. Now, when the platform is allowed to choose the order, the difference is, $\mathcal{R}_{SP}(\{q_1, q_2\}) - \mathcal{R}_{CP}(\{q_2, q_1\}) = \frac{v}{2e}$*.*

B.2 Omitted Proofs

In each subsection we will give full proofs of all results from the main body.

B.2.1 Omitted Proofs from Section 3.2

Proof of Lemma 3. We will derive the desired expressions first for CP, then for SP. First note, CP is equivalent to period L pricing when L goes to T. Therefore,

$$\begin{aligned}\mathcal{R}_{CP}(p) = \lim_{L\to T}\mathcal{R}(L,p) &= \lim_{L\to T}\left(p - \frac{c\left(1 - e^{-qL}\right)}{q}\right)\sum_{i=1}^{\lfloor T/L\rfloor} e^{-(i-1)qL}\overline{\mathrm{F}}(\omega(i)p) \\ &= \left(p - \frac{c\left(1 - e^{-qT}\right)}{q}\right)\overline{\mathrm{F}}\left(\frac{p\left(q - \log(\delta)\right)}{q\left(1 - \delta^T e^{-qT}\right)}\right),\end{aligned}$$

where the third equality follows from the fact that when L goes to T, $\omega(1) = \lim_{L\to T}\frac{q - \log(\delta)}{q(1-\delta^L e^{-qL})} = \frac{q-\log(\delta)}{q(1-\delta^T e^{-qT})}$. This yields Eq. (8), now we will derive Eq. (9). For simplicity of presentation, let $\underline{v} = \frac{p(q-\log(\delta))}{q(1-\delta^T e^{-qT})}$, which is the lowest valuation of a user who will still pay the long term price p. Thus by definition the price satisfies, $p = \int_0^T \underline{v}\delta^t q e^{-qt} dt$. The expected profit earned by CP with price p can be rewritten as,

$$
\begin{aligned}
\mathcal{R}_{CP}(p) &= \left(p - \frac{c\left(1-e^{-qT}\right)}{q}\right) \overline{\mathrm{F}}\left(\frac{p\left(q-\log(\delta)\right)}{q\left(1-\delta^T e^{-qT}\right)}\right) \\
&= \left(p - c\left(\int_0^T tqe^{-qt}dt + Te^{-qT}\right)\right) \overline{\mathrm{F}}(\underline{v}) \\
&= \left(\int_0^T \underline{v}\delta^t q e^{-qt}dt - c\left(-te^{-qt}|_0^T + \int_0^T e^{-qt}dt + Te^{-qT}\right)\right) \overline{\mathrm{F}}(\underline{v}) \\
&= \int_0^T \left(\underline{v}\delta^t q - c\right) e^{-qt}\overline{\mathrm{F}}(\underline{v})dt \\
&= \int_0^T \left(\frac{p\delta^t\left(q-\log(\delta)\right)}{1-\delta^T e^{-qT}} - c\right) e^{-qt}\overline{\mathrm{F}}\left(\frac{p\left(q-\log(\delta)\right)}{q\left(1-\delta^T e^{-qT}\right)}\right) dt,
\end{aligned}
$$

where the first equality is Eq. (8), the second equality follows from rewriting $\frac{c\left(1-e^{-qT}\right)}{q}$ in integral form, the third equality comes from recalling $p = \int_0^T \underline{v}\delta^t q e^{-qt} dt$ by definition of $\underline{v}$ and applying integration by parts, the fourth equality follows by simplifying, and the final equality follows from plugging in via the definition of $\underline{v}$. This yields Eq. (9).

Now we will derive the expressions for SP which is equivalent to period L pricing as L goes to 0. To take this limit, assume the price rate, $p = \frac{p_L}{L}$, of period L pricing is fixed. Then,

$$
\begin{aligned}
\mathcal{R}_{SP}(p) = \lim_{L\to 0}\mathcal{R}(L, p_L) &= \lim_{L\to 0}\left(pL - \frac{c\left(1-e^{-qL}\right)}{q}\right)\sum_{i=1}^{\lfloor T/L\rfloor} e^{-(i-1)qL}\overline{\mathrm{F}}(\omega(i)pL) \\
&= \lim_{L\to 0}\left(\left(p - \frac{c\left(1-e^{-qL}\right)}{qL}\right)\sum_{i=1}^{\lfloor T/L\rfloor} e^{-(i-1)qL}\overline{\mathrm{F}}(\omega(i)pL)\right) L \\
&= \lim_{L\to 0}\left(\left(p - \frac{c\left(1-e^{-qL}\right)}{qL}\right)\sum_{i=1}^{\lfloor T/L\rfloor} e^{-(i-1)qL}\overline{\mathrm{F}}\left(\delta^{-(i-1)L}\frac{\left(q-\log\left(\delta\right)\right)}{q\left(1-\delta^L e^{-qL}\right)}pL\right)\right) L \\
&= \lim_{L\to 0}\left(\left(p - \frac{c\left(1-e^{-qL}\right)}{qL}\right)\sum_{i=1}^{\lfloor T/L\rfloor} e^{-(i-1)qL}\overline{\mathrm{F}}\left(\frac{q-\log\left(\delta\right)}{q}\frac{L}{\left(1-\delta^L e^{-qL}\right)}\delta^{-(i-1)L}p\right)\right) L \\
&= \int_0^T (p - \frac{cq}{q})e^{-qt}\overline{\mathrm{F}}\left(\frac{q-\log\left(\delta\right)}{q}\frac{1}{q-\log(\delta)}\delta^{-t}p\right) dt, \\
&= \int_0^T (p-c)e^{-qt}\overline{\mathrm{F}}(pq^{-1}\delta^{-t})dt,
\end{aligned}
$$

where the first equality follows by definition of period L pricing, and the second equality follows from Eq. (7). The third equality follows from factoring out L. The fourth equality follows by plugging in the definition of $\omega(i)$. The fifth equality follows from rearranging. In the sixth equality we take the limit in L, noting that $\lim_{L\to 0} \frac{L}{1-\delta^L e^{-qL}} = \frac{1}{q-\log(\delta)}$, $\lim_{L\to 0} \frac{1-e^{-qL}}{L} = q$, and $\delta^{-q(i-1)L}$ and $e^{-(i-1)L}$ in the Riemann sum correspond to δ^{-t} and e^{-qt} in the integration. The final equality follows from simplification. This yields Eq. (11), which is an expression of SP as an integration over time t. Finally, we will use Eq. (11) to derive Eq. (10). The profit of SP can be rewritten as,

$$\begin{aligned}
\mathcal{R}_{SP}(p) &= \int_{\frac{p}{q}}^{\infty} (p-c) \left(\int_0^{\min\left\{\frac{\log\left(\frac{p}{vq}\right)}{\log(\delta)}, T\right\}} e^{-qt} dt \right) f(v) dv \\
&= \int_0^T \left(\int_{\frac{p}{q\delta^t}}^{\infty} (p-c)\, f(v) dv \right) e^{-qt} dt \\
&= \int_0^T (p-c)\, e^{-qt} \overline{\mathrm{F}}(pq^{-1}\delta^{-t}) dt,
\end{aligned}$$

where the first equality is a rewriting of Eq. (10), the second equality follows from Fubini's theorem, where $t \le \frac{\log\left(\frac{p}{vq}\right)}{\log(\delta)}$ can be rearranged as $v \ge \frac{p}{q\delta^t}$. The final equality follows from computing the integration over v, from which we obtain Eq. (11). $\square$

Proof of Lemma 4. We will show both the optimal short period price and long period price are unique by studying the first order conditions of their corresponding profit functions.

To demonstrate the uniqueness of the optimal price for SP when F is MHR, first we will show that the optimal price is finite. From Eq. (11), we can upper bound the profit obtained by price p as,

$$\mathcal{R}_{SP}(p) = \int_0^T (p-c)\, \overline{\mathrm{F}}\left(pq^{-1}\delta^{-t}\right) e^{-qt} dt \le \int_0^T (p-c)\, \overline{\mathrm{F}}\left(p\right) e^{-qt} dt,$$

where the inequality follows since $\overline{\mathrm{F}}$ is decreasing. When p tends to infinity, the upper bound on the short period pricing profit goes to 0, which implies the optimal short period price is finite. Now, we show the optimal short period price is unique. To aid our proof, first we will consider general equations of the form:

$$p - c = p\psi(p), \tag{30}$$

where $\psi(p)$ is strictly decreasing and $c \geq 0$. In the following, we first show that the solution of Eq. (30) is unique, then we show the first order condition of $\mathcal{R}_{SP}$ is exactly the form of Eq. (30). To show the uniqueness of the solution to Eq. (30), we will proceed by contradiction. Suppose Eq. (30) has two solutions, denoted by p_1^* and p_2^*, where $p_2^* > p_1^* \geq 0$. By the mean value theorem, there must exists $p \in (p_1^*, p_2^*)$, such that

$$\psi(p) + p\psi'(p) = 1.$$

Thus,

$$\psi(p) = 1 - p\psi'(p) > 1,$$

where the inequality follows by the assumption that $\psi(p)$ is strictly decreasing in p. Also note, by Eq. (30),

$$\psi(p) = \frac{p-c}{p} \leq 1,$$

where the equality follows from rearranging Eq. (30), and the inequality follows from $c \geq 0$, which contradicts with the previous result $\psi(p) > 1$. Therefore, the solution to Eq. (30) is unique. Further, by rearranging Eq. (30), we get

$$p\left(1 - \psi(p)\right) = c,$$

where $p\left(1 - \psi(p)\right)$ is increasing in p since $\psi(p)$ is decreasing in p. Therefore, the solution of Eq. (30) is increasing in c.

Now we will show the first-order condition for the optimal short period price is exactly of the form of Eq. (30) implying uniqueness. Consider the profit function of SP in Eq. (11), its derivative is

$$\frac{\partial \mathcal{R}_{SP}(p)}{\partial p} = \int_0^T \overline{\mathrm{F}}(pq^{-1}\delta^{-t})e^{-qt}dt - (p-c)\int_0^T q^{-1}\delta^{-t} f(pq^{-1}\delta^{-t})e^{-qt}dt = 0. \tag{31}$$

There are two ways the above equation can be zero. The first is if p is such that $\overline{\mathrm{F}}(pq^{-1}) = 0$. In this case, both integrals are zero, and while p is a critical point, such a p cannot be the profit optimal short period price as it earns no profit. Thus assume $\overline{\mathrm{F}}(pq^{-1}) > 0$, then both integrals are positive and we can rearrange the expression Eq. (31) to be

$$
\begin{aligned}
p - c &= \frac{\int_0^T \overline{\mathrm{F}}(pq^{-1}\delta^{-t})e^{-qt}dt}{\int_0^T q^{-1}\delta^{-t} f(pq^{-1}\delta^{-t})e^{-qt}dt} \\
&= \frac{p\log(\delta)\int_0^T \overline{\mathrm{F}}(pq^{-1}\delta^{-t})e^{-qt}dt}{e^{-qT}\overline{\mathrm{F}}(pq^{-1}\delta^{-T}) - \overline{\mathrm{F}}(pq^{-1}) + \int_0^T q\overline{\mathrm{F}}(pq^{-1}\delta^{-t})e^{-qt}dt} \\
&= \frac{p\log(\delta)}{q + \underbrace{e^{-qT}\overline{\mathrm{F}}(pq^{-1}\delta^{-T})\left(\int_0^T \overline{\mathrm{F}}\left(pq^{-1}\delta^{-t}\right)e^{-qt}dt\right)^{-1}}_{g(p)} - \underbrace{\overline{\mathrm{F}}(pq^{-1})\left(\int_0^T \overline{\mathrm{F}}\left(pq^{-1}\delta^{-t}\right)e^{-qt}dt\right)^{-1}}_{h(p)}}
\end{aligned}
$$

where the second equality follows by applying integration by parts to the denominator, and the third equality follows by simplifying. Now to apply Eq. (30), we require that $\frac{\log(\delta)}{q+g(p)-h(p)}$ is decreasing in p (note $\log(\delta) < 0$). To show this, we only need to show $g(p)$ is decreasing in p and $h(p)$ is increasing in p. To show $h(p)$ is increasing in p, consider the derivative of the $h(p)$,

$$
\frac{\partial}{\partial_p} h(p) = \frac{\int_0^T \left(-\overline{\mathrm{F}}(pq^{-1}\delta^{-t})f(pq^{-1}) + \delta^{-t}f(pq^{-1}\delta^{-t})\overline{\mathrm{F}}(pq^{-1})\right)e^{-qt}dt}{q\left(\int_0^T \overline{\mathrm{F}}\left(pq^{-1}\delta^{-t}\right)e^{-qt}dt\right)^2}.
$$

To show that $\frac{\partial}{\partial_p} h(p) \geq 0$, note

$$
\begin{aligned}
\delta^{-t}f(pq^{-1}\delta^{-t})\overline{\mathrm{F}}(pq^{-1}) - \overline{\mathrm{F}}(pq^{-1}\delta^{-t})f(pq^{-1}) &\geq f(pq^{-1}\delta^{-t})\overline{\mathrm{F}}(pq^{-1}) - \overline{\mathrm{F}}(pq^{-1}\delta^{-t})f(pq^{-1}) \\
&= f\left(pq^{-1}\right) f\left(pq^{-1}\delta^{-t}\right)\left(\frac{\overline{\mathrm{F}}(pq^{-1})}{f(pq^{-1})} - \frac{\overline{\mathrm{F}}(pq^{-1}\delta^{-t})}{f(pq^{-1}\delta^{-t})}\right) \\
&\geq 0,
\end{aligned}
$$

where the first inequality follows from $\delta^{-t} \geq 1$ for $t \geq 0$, the equality follows by factoring out $f\left(pq^{-1}\right) f\left(pq^{-1}\delta^{-t}\right)$ and the second inequality from the fact that F is MHR. Thus, $h(p)$ is increasing in p. Next, we show that $g(p)$ is decreasing in p. Consider the derivative of $g(p)$,

$$\frac{\partial}{\partial_p} g(p) = \frac{\int_0^T \left(-\delta^{-T}\overline{\mathrm{F}}(pq^{-1}\delta^{-t}) f(pq^{-1}\delta^{-T}) + \delta^{-t} f(pq^{-1}\delta^{-t})\overline{\mathrm{F}}(pq^{-1}\delta^{-T})\right) e^{-qt} dt}{q\left(\int_0^T \overline{\mathrm{F}}\left(pq^{-1}\delta^{-t}\right) e^{-qt} dt\right)^2}.$$

Like before, to show the derivative of $g(p)$ is negative, consider

$$\begin{aligned}
\delta^{-t} f(pq^{-1}\delta^{-t})\overline{\mathrm{F}}(pq^{-1}\delta^{-T}) - \delta^{-T}\overline{\mathrm{F}}(pq^{-1}\delta^{-t}) f(pq^{-1}\delta^{-T}) &\le \delta^{-T}\left(f(pq^{-1}\delta^{-t})\overline{\mathrm{F}}(pq^{-1}\delta^{-T}) - \overline{\mathrm{F}}(pq^{-1}\delta^{-t}) f(pq^{-1}\delta^{-T})\right) \\
&= \delta^{-T} f\left(pq^{-1}\delta^{-T}\right) f\left(pq^{-1}\delta^{-t}\right)\left(\frac{\overline{\mathrm{F}}(pq^{-1}\delta^{-T})}{f(pq^{-1}\delta^{-T})} - \frac{\overline{\mathrm{F}}(pq^{-1}\delta^{-t})}{f(pq^{-1}\delta^{-t})}\right) \\
&\le 0,
\end{aligned}$$

where the first inequality follows from $\delta^{-T} \ge \delta^{-t}$ for any $t \le T$, the equality follows by factoring out $f\left(pq^{-1}\delta^{-T}\right) f\left(pq^{-1}\delta^{-t}\right)$, and the final inequality from the fact that F is MHR. Thus, $g(p)$ is decreasing in p. Consequently, $\frac{\log(\delta)}{q+g(p)-h(p)}$ is decreasing in p, since $\log(\delta) < 0$. Therefore, the first-order condition of the optimal short period price is exactly of the form in Eq. (30). Combining with the existence of finite optimal short period price, we conclude that the optimal short period price is unique and increasing in c.

The uniqueness of the optimal long period price comes from MHR directly. By taking the derivative of Eq. (8) with respect to p, the first-order condition for the optimal long period price is

$$\overline{\mathrm{F}}\left(\frac{p\left(q - \log(\delta)\right)}{q\left(1 - \delta^T e^{-qT}\right)}\right) - \left(p - \frac{c(1 - e^{-qT})}{q}\right)\left(\frac{q - \log(\delta)}{q\left(1 - \delta^T e^{-qT}\right)}\right) f\left(\frac{p\left(q - \log(\delta)\right)}{q\left(1 - \delta^T e^{-qT}\right)}\right) = 0. \tag{32}$$

As above, assuming p is such that $\overline{\mathrm{F}}\left(\frac{p(q-\log(\delta))}{q(1-\delta^T e^{-qT})}\right) > 0$ (since otherwise p gives no profit and cannot be optimal), we can rearrange Eq. (32) to be

$$p - \frac{c(1 - e^{-qT})}{q} = \left(\frac{q\left(1 - \delta^T e^{-qT}\right)}{q - \log(\delta)}\right) \frac{\overline{\mathrm{F}}\left(\frac{p(q-\log(\delta))}{q(1-\delta^T e^{-qT})}\right)}{f\left(\frac{p(q-\log(\delta))}{q(1-\delta^T e^{-qT})}\right)}.$$

Finally, note $\overline{\mathrm{F}}\left(\frac{p(q-\log(\delta))}{q(1-\delta^T e^{-qT})}\right) f\left(\frac{p(q-\log(\delta))}{q(1-\delta^T e^{-qT})}\right)^{-1}$ is non-increasing in p since F is MHR, while $p - \frac{c(1-e^{-qT})}{q}$ increasing in p. Therefore, we can conclude the optimal long period price is unique and increasing in c. $\square$

B.2.2 Omitted Proofs from Section 3.3

Proof of Lemma 5. Let $p_L^* = \arg\max_p \mathcal{R}(p, L)$ be any optimal price for period L pricing, and consider a feasible short period price $p_S = \frac{qp_L^*}{1-e^{-qL}}$. For any user with valuation v, recall $\tau(v)$ is the largest number of periods the user will pay under period L pricing with price p_L^*. Further, note that $\frac{1-e^{-x}}{x}$ is decreasing in x, therefore,

$$\frac{1-e^{-qL}}{qL} \geq \frac{1-e^{-(q-\log(\delta))L}}{(q-\log(\delta))L}$$

for all $\delta \in (0,1)$. Thus after i periods, for any i less than $\tau(v)$, the user is willing to pay the period L price, and further would also pay the short period price $p = \frac{qp_L^*}{1-e^{-qL}}$ at time iL since the purchasing condition for one implies the other i.e.,

$$v\delta^{iL}\left(\frac{q\left(1-\delta^L e^{-qL}\right)}{q-\log\left(\delta\right)}\right) \geq p_L^* \implies v\delta^{iL} q \geq \frac{qp_L^*}{1-e^{-qL}}.$$

Further, in the i^{th} period, SP with price $p = \frac{qp_L^*}{1-e^{-qL}}$ will yield the same expected profit as subscription pricing of period length L. To show this, note for a user a time t in $[(i-1)L, iL)$ (this corresponds to the i^{th} for period L pricing $i < \tau(v)$), a user will leave only if they match, and the expected profit the user will bring to the platform under SP over the time interval $[(i-1)L, iL)$ is,

$$\int_0^L \left(\frac{qp_L^*}{1-e^{-qL}}\right) e^{-qt} dt = p_L^*.$$

Thus short period pricing with price $p = \frac{qp_L^*}{1-e^{-qL}}$ will yield the same expected profit for all the other periods as under period L pricing, with the exception of the last period the user would pay for. At the moment before the last period starts, the willingness to pay of a user with initial valuation v is $v\delta^{(\tau(v)-1)L}$. To bound the profit earned in this final period, consider a long period pricing with the same price p_L^* but with horizon $T = L$. The profit of this long period pricing with L is at least the profit for the last period of period L pricing, since if a user is willing to pay the last period, they are also willing to pay the long period price p_L^*. Therefore, $\mathcal{R}(L) \leq \mathcal{R}_{SP} + \mathcal{R}_{CP}^L$. □

Proof of Lemma 6. As a preliminary to this proof, we will first introduce one condition on our market parameters that will be helpful.

Condition B.2.1 ((C1) Slow Matching Condition)**.** *When the market parameters q, δ, and T are such that,*

$$\left(1-\frac{q}{\log(\delta)}\right)^{\frac{\log(\delta)}{q}} \geq \left(\frac{q}{q-\log(\delta)}\right)\left(\frac{1-\delta^T e^{-qT}}{1-e^{-qT}}\right),$$

we say the market satisfies a slow matching condition.

To build intuition for (C1), consider the case when the pool of potential matches is large, *i.e.*, $T=\infty$. In this case, the condition reduces to $q \leq -\log(\delta)$.

Now our proof will follow in three steps. First, we reduce the problem to the case where valuations are fixed and deterministic. Next, we bound the ratio for fixed valuations when (C1) holds by analyzing a feasible price, which is the optimal short period price when $c=0$. Finally, we bound the ratio when (C1) does not hold by analyzing a second feasible price, which is the normalized long-period price.

B.2.2.1 Step 1: Reduction to fixed valuations.

Define $\gamma(v) := \min\left\{1-\left(\frac{p}{vq}\right)^{\frac{-q}{\log(\delta)}}, 1-e^{-qT}\right\}$, and note $\gamma(v)$ is a non-decreasing function of v. Now, by Eq. (10) the profit of SP for some fixed price p is,

$$\begin{aligned}\mathcal{R}_{SP}(p) &= \left(\frac{p-c}{q}\right)\int_{p/q}^{\infty}\min\left\{1-\left(\frac{p}{vq}\right)^{\frac{-q}{\log(\delta)}}, 1-e^{-qT}\right\}f(v)dv = \left(\frac{p-c}{q}\right)\int_{p/q}^{\infty}\gamma(v)f(v)dv \\ &= \left(\frac{p-c}{q}\right)\mathbb{E}[\gamma(V)\mathbf{1}_{V\geq p/q}].\end{aligned}$$

By the generalized Markov's inequality, for any $x \geq 0$, we have,

$$\mathbb{E}[\gamma(V)\mathbf{1}_{V\geq p/q}] \geq \overline{\mathrm{F}}(x)\gamma(x),$$

since if $x \geq p/q$, $\gamma(x) \geq 0$, we can apply Markov's inequality, and if $x < p/q$, $\gamma(x) \leq 0$, whereas $\mathbb{E}[\gamma(V)\mathbf{1}_{V\geq p/q}] \geq 0$. Applying the inequality to our expression for $\mathcal{R}_{SP}(p)$ we find,

$$\mathcal{R}_{SP}(p) \geq \left(\frac{p-c}{q}\right)\min\left\{1-\left(\frac{p}{xq}\right)^{\frac{-q}{\log(\delta)}}, 1-e^{-qT}\right\}\overline{\mathrm{F}}(x), \quad \text{for any } x \geq 0. \tag{33}$$

Now, suppose the optimal price for $\mathcal{R}_{CP}(p)$ is p^*, and let $v^* = \frac{p^*(q-\log(\delta))}{q(1-\delta^T e^{-qT})}$. Taking x as v^* in Eq. (33) yields,

$$\mathcal{R}_{SP}(p) \geq \left(\frac{p-c}{q}\right) \min\left\{1 - \left(\frac{p}{v^* q}\right)^{\frac{-q}{\log(\delta)}}, 1 - e^{-qT}\right\} \overline{\mathrm{F}}(v^*). \tag{34}$$

Therefore,

$$\begin{aligned}\frac{\mathcal{R}_{SP}}{\mathcal{R}_{CP}} \geq \frac{\mathcal{R}_{SP}(p)}{\mathcal{R}_{CP}(p^*)} &\geq \frac{\max_p \left(\frac{p-c}{q}\right) \min\left\{1 - \left(\frac{p}{v^* q}\right)^{\frac{-q}{\log(\delta)}}, 1 - e^{-qT}\right\} \overline{\mathrm{F}}(v^*)}{\left(\frac{v^*}{1-\frac{\log(\delta)}{q}}\left(1 - \delta^T e^{-qT}\right) - \frac{c}{q}\left(1 - e^{-qT}\right)\right) \overline{\mathrm{F}}(v^*)} \\ &= \frac{\max_p \left(\frac{p-c}{q}\right) \min\left\{1 - \left(\frac{p}{v^* q}\right)^{\frac{-q}{\log(\delta)}}, 1 - e^{-qT}\right\}}{\left(\frac{v^*}{1-\frac{\log(\delta)}{q}}\left(1 - \delta^T e^{-qT}\right) - \frac{c}{q}\left(1 - e^{-qT}\right)\right)}\end{aligned} \tag{35}$$

where the first inequality follows from using the optimal price for CP, and the second inequality follows by applying Eq. (34) and plugging in the profit for CP from Eq. (8). Define F_v to be the point mass distribution for a random variable that is equal to some constant v with probability one. Note, then that the ratio in Eq. (35) is exactly the same as the ratio for F_{v^*}, i.e.:

$$\frac{\max_p \left(\frac{p-c}{q}\right) \min\left\{1 - \left(\frac{p}{v^* q}\right)^{\frac{-q}{\log(\delta)}}, 1 - e^{-qT}\right\}}{\left(\frac{v^*}{1-\frac{\log(\delta)}{q}}\left(1 - \delta^T e^{-qT}\right) - \frac{c}{q}\left(1 - e^{-qT}\right)\right)} \geq \inf_v \frac{\mathcal{R}_{SP}^v}{\mathcal{R}_{CP}^v}. \tag{36}$$

where we use $\mathcal{R}_{SP}^v$ and $\mathcal{R}_{CP}^v$ to donote the profit of SP and CP with point mass valuations v, respectively. For the remainder of our proof, we will lower bound Eq. (36) by finding the worst case ratio over all point mass valuations v.

B.2.2.2 Step 2: Bounds when (C1) holds.

Fix some point mass valuation v. First assume, that $\delta^T > \left(1-\frac{q}{\log(\delta)}\right)^{\frac{\log(\delta)}{q}}$, and consider the feasible short period price be $\tilde{p} = vq\delta^T$. The corresponding short period profit is

$$\begin{aligned}
\mathcal{R}^v_{SP}(\tilde{p}) &= \left(v\delta^T - \frac{c}{q}\right)\left(1-e^{-qT}\right) \\
&= \mathcal{R}^v_{CP}\left(\frac{\delta^T\left(1-e^{-qT}\right)}{\frac{q}{q-\log(\delta)}\left(1-\delta^T e^{-qT}\right)}\right) + \frac{c\left(1-e^{-qT}\right)}{q}\left[\frac{\delta^T\left(1-e^{-qT}\right)}{\frac{q}{q-\log(\delta)}\left(1-\delta^T e^{-qT}\right)} - 1\right].
\end{aligned}$$

Rearranging we have,

$$\begin{aligned}
\frac{\mathcal{R}^v_{SP}(\tilde{p})}{\mathcal{R}^v_{CP}(p^*)} &= \frac{\delta^T\left(1-e^{-qT}\right)}{\frac{q}{q-\log(\delta)}\left(1-\delta^T e^{-qT}\right)} + \frac{c\left(1-e^{-qT}\right)}{q}\left[\frac{\delta^T\left(1-e^{-qT}\right)}{\frac{q}{q-\log(\delta)}\left(1-\delta^T e^{-qT}\right)} - 1\right] \\
&\geq \left(1-\frac{q}{\log(\delta)}\right)^{\frac{\log(\delta)}{q}}\left(\frac{\left(1-e^{-qT}\right)}{\frac{q}{q-\log(\delta)}\left(1-\delta^T e^{-qT}\right)}\right) \\
&\quad + \frac{c\left(1-e^{-qT}\right)}{q}\left[\left(1-\frac{q}{\log(\delta)}\right)^{\frac{\log(\delta)}{q}}\left(\frac{\left(1-e^{-qT}\right)}{\frac{q}{q-\log(\delta)}\left(1-\delta^T e^{-qT}\right)}\right) - 1\right] \\
&\geq 1 + \frac{c\left(1-e^{-qT}\right)}{q}\left[1-1\right] = 1.
\end{aligned}$$

where the first inequality follows from our assumption, and the second inequality follows from rearranging (C1).

Next consider the alternative assumption, $\delta^T \leq \left(1-\frac{q}{\log(\delta)}\right)^{\frac{\log(\delta)}{q}}$. For this case, consider the feasible short period price $\tilde{p} = vq\left(\frac{\log(\delta)-q}{\log(\delta)}\right)^{\frac{\log(\delta)}{q}}$. By Eq. (10), the profit generated by $\tilde{p}$ is,

$$\begin{aligned}
\mathcal{R}^v_{SP}(\tilde{p}) &= \frac{vq}{q-\log(\delta)}\left(1-\frac{q}{\log(\delta)}\right)^{\frac{\log(\delta)}{q}} - \frac{c}{q-\log(\delta)} \\
&= \frac{\mathcal{R}^v_{CP}}{1-\delta^T e^{-qT}}\left(1-\frac{q}{\log(\delta)}\right)^{\frac{\log(\delta)}{q}} + \frac{c}{q}\left[\left(1-\frac{q}{\log(\delta)}\right)^{\frac{\log(\delta)}{q}}\left(\frac{1-e^{-qT}}{1-\delta^T e^{-qT}}\right) - \frac{q}{q-\log(\delta)}\right].
\end{aligned}$$

By (C1), we always have

$$\left(1-\frac{q}{\log(\delta)}\right)^{\frac{\log(\delta)}{q}}\left(\frac{1-e^{-qT}}{1-\delta^T e^{-qT}}\right) - \frac{q}{q-\log(\delta)} \geq 0.$$

Therefore,

$$\mathcal{R}_{SP}^{v} \geq \mathcal{R}_{CP}^{v} \left(1 - \frac{q}{\log(\delta)}\right)^{\frac{\log(\delta)}{q}} \left(1 - \delta^{T} e^{-qT}\right)^{-1},$$

as desired. Finally, letting $x = \frac{q}{-\log(\delta)}$, and minimizing the above expression for $0 < x < 1$, we have $\min_{x \in (0,1)} (1+x)^{\frac{-1}{x}} \geq \frac{1}{e}$ with the minimum occurring as x tends to 0. Thus in this case $\frac{\mathcal{R}_{SP}}{\mathcal{R}_{CP}} \geq \frac{1}{e}$.

B.2.2.3 Step 3: Bounds when (C1) does not hold.

Again fix some point mass valuation v, and now consider the feasible short period price, $\tilde{p} = vq^2(1 - \delta^T e^{-qT}) \left((q - \log(\delta)) \left(1 - e^{-qT}\right)\right)^{-1}$. By Eq. (10), the profit generated by $\tilde{p}$ is,

$$\begin{aligned}
\mathcal{R}_{SP}^{v}(\tilde{p}) &= \left(\frac{vq\left(1 - \delta^T e^{-qT}\right)}{(q - \log(\delta))\left(1 - e^{-qT}\right)} - \frac{c}{q}\right) \left[1 - \left(\left(\frac{1}{1 - \frac{\log(\delta)}{q}}\right)\left(\frac{1 - \delta^T e^{-qT}}{1 - e^{-qT}}\right)\right)^{-\frac{q}{\log(\delta)}}\right] \\
&= \frac{\mathcal{R}_{CP}^{v}}{1 - e^{-qT}} \left[1 - \left(\left(\frac{1}{1 - \frac{\log(\delta)}{q}}\right)\left(\frac{1 - \delta^T e^{-qT}}{1 - e^{-qT}}\right)\right)^{-\frac{q}{\log(\delta)}}\right]
\end{aligned}$$

Rearranging we have,

$$\frac{\mathcal{R}_{SP}^{v}(\tilde{p})}{\mathcal{R}_{CP}^{v}} \geq \underbrace{\frac{1}{1 - e^{-qT}} \left[1 - \left(\left(\frac{1}{1 - \frac{\log(\delta)}{q}}\right)\left(\frac{1 - \delta^T e^{-qT}}{1 - e^{-qT}}\right)\right)^{-\frac{q}{\log(\delta)}}\right]}_{\phi(T)}.$$

The derivative of $\phi(T)$ with respect to T is,

$$\frac{\partial \phi(T)}{\partial T} = \left(\frac{qe^{-qT}}{\left(1 - e^{-qT}\right)^2}\right) \left(\left(1 - \frac{q}{\log(\delta)}\right)\left(\frac{1 - \delta^T}{1 - \delta^T e^{-qT}}\right)\left(\frac{1 - \delta^T e^{-qT}}{\left(1 - \frac{\log(\delta)}{q}\right)\left(1 - e^{-qT}\right)}\right)^{-\frac{q}{\log(\delta)}} - 1\right).$$

If $\frac{q}{-\log(\delta)} \geq 1$, then $\frac{\partial \phi(T)}{\partial T} \geq 0$, and thus,

$$\frac{\mathcal{R}_{SP}^{v}(\tilde{p})}{\mathcal{R}_{CP}^{v}} \geq \lim_{T \to 0} \phi(T) = \frac{1}{2}.$$

If $\frac{q}{-\log(\delta)} < 1$, then $\frac{\partial \phi(T)}{\partial T} < 0$. Define T^* such that,

$$\left(1 - \frac{q}{\log(\delta)}\right)^{\frac{\log(\delta)}{q}} = \left(\frac{q}{q - \log(\delta)}\right)\left(\frac{1 - \delta^{T^*} e^{-qT^*}}{1 - e^{-qT^*}}\right),$$

in this case, (C1) does not hold only for $T < T^*$. Thus to complete the proof we can consider

$$\frac{\mathcal{R}^v_{SP}(\tilde{p})}{\mathcal{R}^v_{CP}} \geq \lim_{T \to T^*} \phi(T) = \left(1 - \frac{q}{\log(\delta)}\right)^{\frac{\log(\delta)}{q}} \left(1 - \delta^{T^*} e^{-qT^*}\right)^{-1}.$$

Combining across all cases and taking the minimum yields the claimed bound.

Finally, for tightness, consider the case when $c = 0$, $T = \infty$, and valuations are a point mass v. In this case the optimal price for SP can be computed as $p^* = vq \left(\frac{\log(\delta) - q}{\log(\delta)}\right)^{\frac{\log(\delta)}{q}}$ yielding optimal profit

$$\mathcal{R}^v_{SP}(0) = \frac{vq}{q - \log(\delta)} \left(1 - \frac{q}{\log(\delta)}\right)^{\frac{\log(\delta)}{q}}.$$

Then the ratio between SP and CP is

$$\frac{\mathcal{R}^v_{SP}(0)}{\mathcal{R}^v_{CP}(0)} = \left(1 - \frac{q}{\log(\delta)}\right)^{\frac{\log(\delta)}{q}}$$

matching the guarantee. Taking $\frac{q}{-\log(\delta)} \to 0$ gives the $1/e$ constant factor. □

Proof of Theorem 6. The proof will follow by applying Lemmas 5 and 6. First, in Lemma 5 we show that the optimal profit of period L pricing can be upper bounded by the profit of SP plus the profit of CP where the time to exhaustion T equals L. Further, note that by looking at the integration formulation of SP in Eq. (11), the profit of SP is monotonically increasing in T i.e., $\mathcal{R}^{T_1}_{SP} \leq \mathcal{R}^{T_2}_{SP}$ for any $T_1 \leq T_2$, where we use $\mathcal{R}^T_{SP}$ to denote profit of SP with T. Therefore, by applying Lemma 6 as follows,

$$\mathcal{R}(L) \leq \mathcal{R}_{SP} + \mathcal{R}^L_{CP} \leq \mathcal{R}_{SP} + e\mathcal{R}^L_{SP} \leq (1 + e)\mathcal{R}_{SP},$$

we obtain the desired result. Finally, we construct a family of examples in Example 6 to show that no other fixed choice of pricing strategy besides short period pricing can guarantee any constant factor of the optimal profit. □

Proof of Theorem 7. First note that when c = 0, the optimal long period pricing profit is

$$\mathcal{R}_{CP} = \max_p p\overline{\mathrm{F}}\left(\frac{q-\log(\delta)}{q\left(1-\delta^T e^{-qT}\right)}p\right) = \frac{q\left(1-\delta^T e^{-qT}\right)}{q-\log(\delta)}p^*\overline{\mathrm{F}}(p^*),$$

where p^* is the solution of $\max_p p\overline{\mathrm{F}}(p)$. Now, using Eq. (7), we can upper bound $\mathcal{R}(L)$ as,

$$\begin{aligned}
\mathcal{R}(L) &= \max_p \sum_{i=1}^{\lfloor T/L \rfloor} pe^{-(i-1)qL}\overline{\mathrm{F}}\left(\omega(i)p\right) \\
&\leq \sum_{i=1}^{\lfloor T/L \rfloor} \max_p pe^{-(i-1)qL}\overline{\mathrm{F}}\left(\omega(i)p\right) \\
&= \sum_{i=1}^{\lfloor T/L \rfloor} \frac{p^*}{\omega(i)}\overline{\mathrm{F}}\left(p^*\right)e^{-(i-1)qL} \\
&= p^*\overline{\mathrm{F}}\left(p^*\right)\sum_{i=1}^{\lfloor T/L \rfloor} \frac{e^{-(i-1)qL}}{\omega(i)} \\
&= p^*\overline{\mathrm{F}}\left(p^*\right)\sum_{i=1}^{\lfloor T/L \rfloor} e^{-(i-1)qL}\left(\frac{q\delta^{(i-1)L}\left(1-\delta^L e^{-qL}\right)}{q-\log(\delta)}\right) \\
&= p^*\overline{\mathrm{F}}\left(p^*\right)\left(\frac{q\left(1-\delta^L e^{-qL}\right)}{q-\log(\delta)}\right)\sum_{i=1}^{\lfloor T/L \rfloor} e^{-(i-1)qL}\delta^{(i-1)L} \\
&= p^*\overline{\mathrm{F}}\left(p^*\right)\left(\frac{q\left(1-\delta^L e^{-qL}\right)}{q-\log(\delta)}\right)\left(\frac{1-\delta^{L\lfloor T/L \rfloor}e^{-qL\lfloor T/L \rfloor}}{1-\delta^L e^{-qL}}\right) \\
&= p^*\overline{\mathrm{F}}\left(p^*\right)\left(\frac{q}{q-\log(\delta)}\right)\left(1-\delta^{L\lfloor T/L \rfloor}e^{-qL\lfloor T/L \rfloor}\right) \\
&\leq \mathcal{R}_{CP},
\end{aligned}$$

the first inequality follows by exchanging the max and sum, the second equality follows by taking the maximum, the third and fourth equality comes from the definition of $\omega(i)$ in Eq. (5), the final inequality follows $1-\delta^{L\lfloor T/L \rfloor}e^{-qL\lfloor T/L \rfloor} \leq 1-\delta^T e^{-qT}$.

□

B.2.3 Omitted Proofs of Section 3.4

Proof of Lemma 7. Recall the profit expression of period L pricing for price p in Eq. (7),

$$\begin{aligned}
\mathcal{R}(p,L) &= \left(p - \frac{c\left(1-e^{-qL}\right)}{q}\right) \sum_{i=1}^{\lfloor T/L \rfloor} \left(\frac{1-e^{-iqL}}{1-e^{-qL}}\right) \left(\overline{\mathrm{F}}(\omega(i)p) - \overline{\mathrm{F}}(\omega(i+1)p)\right) \\
&\quad + \left(p - \frac{c\left(1-e^{-qL}\right)}{q}\right) \left(\frac{1-e^{-\lfloor T/L \rfloor qL}}{1-e^{-qL}}\right) \overline{\mathrm{F}}\left(\omega\left(\lfloor T/L \rfloor + 1\right) p\right) \\
&= \left(\frac{p}{1-e^{-qL}} - \frac{c\left(1-e^{-qL}\right)}{q}\right) \sum_{i=1}^{\lfloor T/L \rfloor} \left(1-e^{-iqL}\right) \left(\overline{\mathrm{F}}(\omega(i)p) - \overline{\mathrm{F}}(\omega(i+1)p)\right) \\
&\quad + \left(\frac{p}{1-e^{-qL}} - \frac{c\left(1-e^{-qL}\right)}{q}\right) \left(1-e^{-\lfloor T/L \rfloor qL}\right) \overline{\mathrm{F}}\left(\omega\left(\lfloor T/L \rfloor + 1\right) p\right) \\
&= \left(\frac{p}{1-e^{-qL}} - \frac{c}{q}\right) \mathcal{M}(p,L),
\end{aligned}$$

where the final equality follows from the definition of match proportion $\mathcal{M}(p,L)$.

□

Proof of Theorem 8. Starting from Eq. (11), note that when $c=0$ and $T=\infty$, the profit of SP from any price p simplifies to,

$$\mathcal{R}_{SP}(p) = \int_0^\infty p\overline{\mathrm{F}}(pq^{-1}\delta^{-t})e^{-qt}dt = \frac{p}{q}\mathcal{M}_{SP}(p).$$

where the final equality follows from Lemma 7. Therefore,

$$\mathcal{M}_{SP}(p) = \int_0^\infty q\overline{\mathrm{F}}(pq^{-1}\delta^{-t})e^{-qt}dt. \tag{37}$$

Now, by taking the derivative of the profit with respect to p, the first order condition for SP in this case is,

$$\frac{\partial \mathcal{R}_{SP}(p)}{\partial p} = \int_0^\infty \overline{\mathrm{F}}(pq^{-1}\delta^{-t})e^{-qt}dt - p\int_0^\infty q^{-1}\delta^{-t}f(pq^{-1}\delta^{-t})e^{-qt}dt = 0. \tag{38}$$

Let p_S^* denote the unique optimal price for SP. It satisfies the first order condition above, and further

$$\begin{aligned}
\mathcal{M}_{SP}(p_S^*) &= \int_0^\infty q\overline{\mathrm{F}}(p_S^* q^{-1}\delta^{-t})e^{-qt}dt && \text{Eq. (37)}\\
&= \int_0^\infty p_S^*\delta^{-t} f(p_S^* q^{-1}\delta^{-t})e^{-qt}dt && \text{Rearranging Eq. (38)}\\
&= \frac{q}{-\log(\delta)}\overline{\mathrm{F}}(p_S^* q^{-1}) + \int_0^\infty \frac{q}{\log(\delta)} qe^{-qt}\overline{\mathrm{F}}(p_S^* q^{-1}\delta^{-t})dt && \text{Integration by parts}\\
&= \frac{q}{-\log(\delta)}\overline{\mathrm{F}}(p_S^* q^{-1}) + \frac{q}{\log(\delta)}\mathcal{M}_{SP}(p_S^*) && \text{Eq. (37) again,}\\
\implies \mathcal{M}_{SP}(p_S^*) &= \frac{q}{q-\log(\delta)}\overline{\mathrm{F}}(p_S^* q^{-1}) && \text{Rearranging}
\end{aligned}$$

Armed with this expression for the match proportion of SP under the profit maximizing price, now consider the feasible long period price $p_C = \frac{p_S^*}{q-\log(\delta)}$. By Eq. (8) and Lemma 7, the match proportion of CP under this price will be,

$$\mathcal{M}_{CP}(p_C) = \overline{\mathrm{F}}\left(\frac{p_C\left(q-\log(\delta)\right)}{q}\right) = \overline{\mathrm{F}}(p_S^* q^{-1}) = \left(\frac{q-\log(\delta)}{q}\right)\mathcal{M}_{SP}.$$

This implies the desired match proportion guarantee, to show simultaneous profit dominance of this feasible price note by Eq. (8),

$$\mathcal{R}_{CP}(p_C) = p_C\overline{\mathrm{F}}\left(\frac{p_C\left(q-\log(\delta)\right)}{q}\right) = \frac{p_S^*}{q-\log(\delta)}\overline{\mathrm{F}}(p_S^* q^{-1}) = \frac{p_S^*}{q}\mathcal{M}_{SP} = \mathcal{R}_{SP}.$$

Thus under the price p_C, CP earns the same profit as $\mathcal{R}_{SP}$, while matching more users.

$\square$

B.2.4 Omitted Proofs from Section 3.5

Proof. Proof of Theorem 9. We will prove the three parts separately. First, we show the result for two match rates $\{q_1, q_2\}$ with population sizes $\{t_1, t_2\}$. Let $q = \frac{q_1 t_1 + q_2 t_2}{t_1 + t_2}$ be the average match rate and recall the profit of CP with long period price p is,

$$\mathcal{R}_{CP}(p, \sigma(q)) = (p - c\mathbb{E}[\text{Time on platform}])\,\overline{\mathrm{F}}\left(\left(\frac{q - \log(\delta)}{q - q\delta^T e^{-qT}}\right) p\right).$$

For any price p, the proportion of users which will pay the price depends only on the user's supposed match rate q and is thus independent of the order in which candidate matches are shown. Thus, to compare the profit of the two possible orderings, $\{q_2, q_1\}$ and $\{q_1, q_2\}$, we only need to compare the expected matching cost to the platform from each user, or equivalently, the expected time on the platform. Let $X_{CP}(\sigma(q))$ be the expected time a user who pays the long period price will spend on the platform under match rate order $\sigma(q)$. When the platform chooses matching order $\{q_1, q_2\}$, the expected time on the platform for a user who paid the long period price is,

$$\begin{aligned}
X_{CP}(\{q_1, q_2\}) = \mathbb{E}[\text{Time on platform}] &= \int_0^{t_1} tq_1 e^{-tq_1} dt + e^{-t_1 q_1} \int_0^{t_2} (t + t_1)\, q_2 e^{-tq_2} dt + Te^{-(t_1 q_1 + t_2 q_2)} \\
&= \int_0^{t_1} tq_1 e^{-tq_1} dt + e^{-t_1 q_1} \left(\int_0^{t_2} tq_2 e^{-tq_2} dt + t_1 \left(1 - e^{-q_2 t_2}\right)\right) + Te^{-(t_1 q_1 + t_2 q_2)}.
\end{aligned}$$

Similarly, when the matching order is $\{q_2, q_1\}$, the expected time on the platform for a user who paid the long period price is

$$\begin{aligned}
X_{CP}(\{q_2, q_1\}) = \mathbb{E}[\text{Time on platform}] &= \int_0^{t_2} tq_2 e^{-tq_2} dt + e^{-t_2 q_2} \int_0^{t_1} (t + t_2)\, q_1 e^{-tq_1} dt + Te^{-(t_1 q_1 + t_2 q_2)} \\
&= \int_0^{t_2} tq_2 e^{-tq_2} dt + e^{-t_2 q_2} \left(\int_0^{t_1} tq_1 e^{-tq_1} dt + t_2 \left(1 - e^{-t_2 q_2}\right)\right) + Te^{-(t_1 q_1 + t_1 q_1)}.
\end{aligned}$$

Then,

$$\begin{aligned}
&X_{CP}(\{q_1, q_2\}) - X_{CP}(\{q_2, q_1\}) \\
&= \left(1 - e^{-t_2 q_2}\right)\left(\int_0^{t_1} tq_1 e^{-tq_1} dt + t_1 e^{-t_1 q_1}\right) - \left(1 - e^{-t_1 q_1}\right)\left(\int_0^{t_2} tq_2 e^{-tq_2} dt + t_2 e^{-t_2 q_2}\right), \\
&= \frac{\left(1 - e^{-t_1 q_1}\right)\left(1 - e^{-t_2 q_2}\right)}{q_1} - \frac{\left(1 - e^{-t_1 q_1}\right)\left(1 - e^{-t_2 q_2}\right)}{q_2} \geq 0,
\end{aligned}$$

where the final inequality follows from $q_2 \geq q_1$. Therefore, we conclude for any distribution F, $\mathcal{R}_{CP}(\{q_2, q_1\}) \geq \mathcal{R}_{CP}(\{q_1, q_2\})$.

Now, for k match rates $\{q_1, \ldots, q_k\}$, with associated populations $\{t_1, \ldots, t_k\}$, let $T_i = \sum_{j=1}^{i} t_j$, for $i = 1, \ldots k$, and $T_0 = 0$. The expected time on the platform for a user who paid the long period price is

$$X_{CP}(\{q_1, \ldots, q_k\}) = \sum_{i=1}^{k} \left(\int_0^{t_i} (t + T_{i-1})\, e^{-q_i(t+T_{i-1})} dt \right) + T e^{-qT}.$$

If we swap q_j and q_{j+1} in $\{q_1, \ldots, q_k\}$ where $1 \leq j \leq k-1$, the expected time on the platform for a user who paid the long period price becomes

$$\begin{aligned}
&X_{CP}(\{q_1, \ldots q_{j+1}, q_j, \ldots, q_k\}) \\
&= T e^{-qT} + \sum_{i=1}^{j-1} \left(\int_0^{t_i} (t + T_{i-1})\, e^{-q_i(t+T_{i-1})} dt \right) + \int_0^{t_{j+1}} (t + T_{j-1}) e^{-q_{j+}(t+T_{j-1})} dt \\
&+ \int_0^{t_j} (t + T_{j-1} + t_{j+1}) e^{-q_{j+1}(t+T_{j-1}+t_{j+1})} dt + \sum_{i=j+2}^{k} \left(\int_0^{t_i} (t + T_{i-1})\, e^{-q_i(t+T_{i-1})} dt \right)
\end{aligned}$$

Note all the other parts for the integrations do not change if we only swap two match rates next to each other. Therefore, we can generalize the proof for order $\{q_1, \ldots, q_k\}$ by switching any two reverse orders that are next to each other.

As in part a), we first show the result for two match rates $\{q_1, q_2\}$ with population sizes $\{t_1, t_2\}$. Again, let $q = \frac{q_1 t_1 + q_2 t_2}{t_1 + t_2}$ be the average match rate and recall profit of SP with short period price p is,

$$\mathcal{R}_{SP}(p\sigma(q)) = \int_{\frac{p}{q}}^{\infty} (p - c) \times \mathbb{E}\left[\text{Time on platform}|\text{SP}\right] f(v) dv.$$

The profit dominance $\mathcal{R}_{SP}(p, \{q_1, ..., q_K\}) \geq \mathcal{R}_{SP}(p, \sigma(q))$ will follow if we can show that for each user, the expected time on the platform is be longer under the match rate $\{q_1, ..., q_K\}$. Let $X_{SP}(v, p, \{q_1, q_2\})$ be the expected time a user with valuation v will stay on the platform under price p, we only need to show that $X_{SP}(v, p, \{q_1, q_2\})$ satisfies

$$X_{SP}(v, p, \{q_1, q_2\}) \geq X_{SP}(v, p, \{q_2, q_1\}).$$

If T is binding, the expected time on platform will be the same as long period pricing model, therefore the conclusion is the same. Otherwise, let $\mu(v) = \frac{\log(p/vq)}{\log(\delta)}$, when $\mu(v) \leq \min\{t_1, t_2\}$,

$$X_{SP}\left(v, p, \{q_1, q_2\}\right) = \int_0^{\mu(v)} e^{-tq_1} dt, \quad X_{SP}\left(v, p, \{q_2, q_1\}\right) = \int_0^{\mu(v)} e^{-tq_2} dt.$$

Note that $e^{-tq_1} \geq e^{-tq_2}$, therefore,

$$X_{SP}(v, p, \{q_1, q_2\}) \geq X_{SP}(v, p, \{q_2, q_1\}).$$

When $t_1 \leq \mu(v) \leq t_2$

$$\begin{aligned} X_{SP}(v, p, \{q_1, q_2\}) &= \int_0^{t_1} e^{-tq_1} dt + e^{-qt_1} \int_0^{\mu(v)-t_1} e^{-tq_2} dt, \\ X_{SP}(v, p, \{q_2, q_1\}) &= \int_0^{t_1} e^{-tq_2} dt + e^{-qt_2} \int_0^{\mu(v)-t_1} e^{-tq_2} dt. \end{aligned}$$

Similarly, we can show that

$$X_{SP}(v, p, \{q_1, q_2\}) \geq X_{SP}(v, p, \{q_2, q_1\})$$

for $t_2 \leq \mu(v) \leq t_2$ or $\max\{t_1, t_2\} \leq \mu(v) \leq t_1 + t_2$. Therefore, we can conclude that for any distribution F, $\mathcal{R}_{SP}(p, \{q_2, q_1\}) \leq \mathcal{R}_{SP}(p, \{q_1, q_2\})$.

For matching order $\{q_1, \ldots, q_k\}$, let $T_i = \sum_{j=1}^{i} t_j$, for $i = 1, \ldots k$, and $T_0 = 0$. If T is binding, the expected time on the platform for a user whose valuation is v under the short period price p is

$$X_{SP}(v, p, \{q_1, \ldots, q_k\}) = \sum_{i=1}^{k} \left(\int_0^{t_i} \left(t + T_{i-1}\right) e^{-q_i(t+T_{i-1})} dt \right) + Te^{-qT},$$

and we can apply the proof in part a). Otherwise, let $q_{k'}$ be the k'-th match rate such that $T_{k'} \leq \mu(v) \leq T_{k'+1}$, the expected time on the platform for a user whose valuation is v under the short period price p is

$$\begin{aligned} X_{SP}(v, p, \{q_1, \ldots, q_k\}) = &\sum_{i=1}^{k'} \left(\int_0^{t_i} \left(t + T_{i-1}\right) e^{q_i(t+T_{i-1})} dt \right) \\ &+ \int_0^{\mu(v)-T_{k'}} \left(t + T_{k'}\right) e^{q_i(t+T_{k'})} dt + \mu(v) e^{-\left(\sum_{i=1}^{k'} q_i t_i + q_{k'+1}(\mu(v)-T_{k'})\right)}. \end{aligned}$$

If we swap match rates q_j and q_{j+1} where $j < k'$, it will be the same as T is binding, if we swap match rates q_j and q_{j+1} where $j \geq k'+1$, it will cause no difference for user's the expected time on the platform, if we swap $q_{k'}$ and $q_{k'+1}$, the expected before k' will stay as $\sum_{i=1}^{k'-1}\left(\int_0^{t_i}(t+T_{i-1})\,e^{q_i(t+T_{i-1})}dt\right)$, we only need to consider the difference in $q_{k'}$ and $q_{k'+1}$, which is analyzed above. Therefore, we can generalize the proof for order $\{q_1, ..., q_k\}$ by switching any two reverse orders that next to each other.

As in Theorem 9(a), let $X_{CP}(\sigma(q))$ be the expected time a user who pays the long period price will spend on the platform under match rate order $\sigma(q)$, and let $X_{CP}(\{q_k, \ldots, q_1\})$ be the expected time a user who pays the long period price will spend on the platform when the match rate order is $\{q_k, \ldots, q_1\}$. First we will show $X_{CP}(\{q_k, \ldots, q_1\}) \leq X_{CP}(\sigma(q))$, i.e., compared with any other order $\sigma(q)$, users will leave the platform sooner when match rates are in descending order. To this end, the profit of long period pricing is

$$\mathcal{R}_{CP}(p, \{q_k, \ldots, q_1\}) = (p - cX_{CP}(\{q_k, \ldots, q_1\}))\,\overline{\mathrm{F}}\left(\frac{q-\log(\delta)}{q\left(1-\delta^T e^{-qT}\right)}p\right).$$

Let p^* be the optimal long period price when the match rate order is $\sigma(q)$, i.e. p^* such that,

$$\frac{\partial\mathcal{R}_{CP}(p,\sigma(q))}{\partial p}|_{p^*} = \overline{\mathrm{F}}\left(\frac{q-\log(\delta)}{q\left(1-\delta^T e^{-qT}\right)}p^*\right) - \left(\frac{q-\log(\delta)}{q\left(1-\delta^T e^{-qT}\right)}\right)(p^* - cX_{CP}(\sigma(q)))\,f\left(\frac{q-\log(\delta)}{q\left(1-\delta^T e^{-qT}\right)}p^*\right) = 0.$$

Let $v^* = \frac{q-\log(\delta)}{q(1-\delta^T e^{-qT})}p^*$, then, the first order derivative of $\mathcal{R}_{CP}(p, \{q_k, \ldots, q_1\})$ at p^* is

$$\begin{aligned}
&\frac{\partial\mathcal{R}_{CP}(p,\{q_k,\ldots,q_1\})}{\partial p}|_{p^*} = \overline{\mathrm{F}}\left(v^*\right) - \left(\frac{q-\log(\delta)}{q\left(1-\delta^T e^{-qT}\right)}\right)(p^* - cX_{CP}(\{q_k,\ldots,q_1\}))\,f\left(v^*\right) \\
&= \frac{\partial\mathcal{R}_{CP}(p,\sigma(q))}{\partial p}|_{p^*} + c\left(\frac{q-\log(\delta)}{q\left(1-\delta^T e^{-qT}\right)}\right)(X_{CP}(\{q_k,\ldots,q_1\} - X_{CP}(\sigma(q)))\,f(v^*) \\
&= c\left(\frac{q-\log(\delta)}{q\left(1-\delta^T e^{-qT}\right)}\right)(X_{CP}(\{q_k,\ldots,q_1\} - X_{CP}(\sigma(q)))\,f(v^*) \leq 0,
\end{aligned}$$

where the second equality follows from minus $X_{CP}(\sigma(q))$, then add $X_{CP}(\sigma(q))$ back, the third equality follows from $\frac{\partial\mathcal{R}_{CP}(p,c,F,\{T\},\sigma(q))}{\partial p}|_{p^*} = 0$, the inequality follows from $X_{CP}(\{q_k, \ldots, q_1\}) \leq X_{CP}(\sigma(q))$, which is shown in Theorem 9(a). Thus to maximize the profit, the optimal long period price for $\mathcal{R}_{CP}(\{q_k, \ldots, q_1\})$ should be lower than p^*.

$\square$

B.3 Additional Figures

Figure 24: Examples of subscription and long period pricing in online dating platforms.

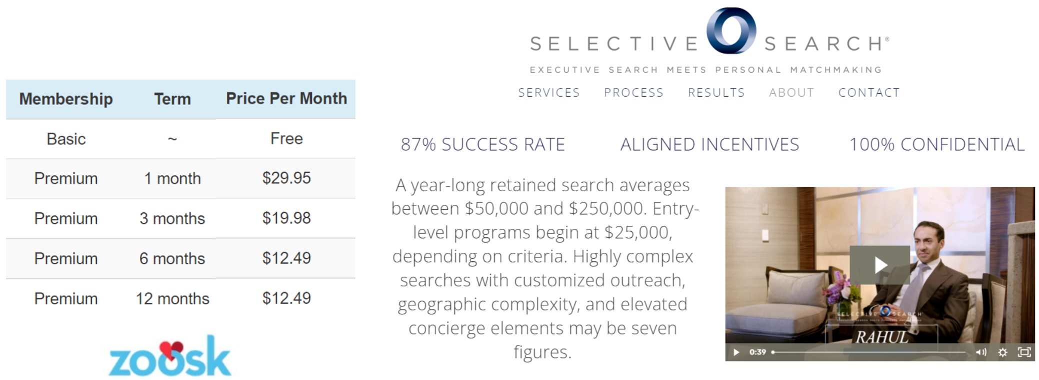

Note. Depicted are price offerings for two online dating platforms. On the left is an example of short period price at `https://www.zoosk.com/`. On the right is an example of long period pricing at `https://www.selectivesearch.com/pricing`.

Figure 25: Relations between optimal price, profit, and match proportion when valuations are uniform.

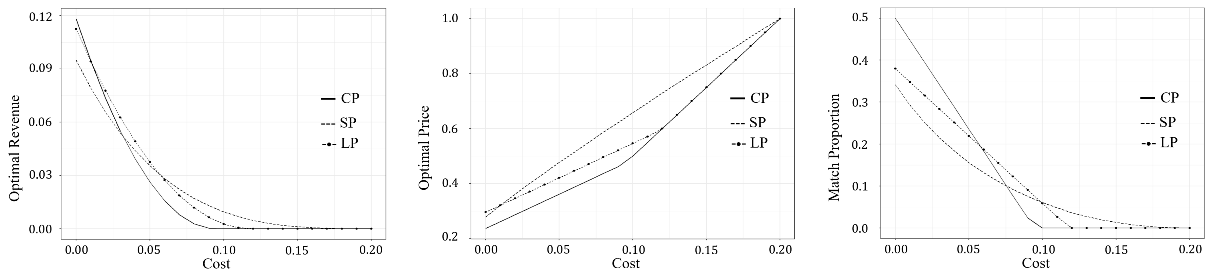

Note. Here we plot the profit, optimal prices, and match proportions under SP, CP and period L pricing with $L = T/7$, when valuations are drawn from an uniform(0, 1) distribution, and where $T = 50$, $\delta = 0.8$, $q = 0.2$, and c varies. In the left panel, we plot the profits of SP, CP and period L pricing, and the note relative profit ordering switches from $\mathcal{R}_{CP} > \mathcal{R}_{SP}$ when $c \leq 0.03$, to $\mathcal{R}_{CP} > \mathcal{R}_{SP}$ for $c > 0.03$. In the middle panel, we plot the normalized optimal long period, short period and period L prices. In the right panel, we plot the proportion of the market that gets matched under SP, CP and period L pricing. Note that $\mathcal{M}_{CP}$ dominates $\mathcal{M}_{SP}$ for $c \leq 0.07$.

Figure 26: Relations between optimal price, profit, and match proportion when valuations are a mixture of log-normal.

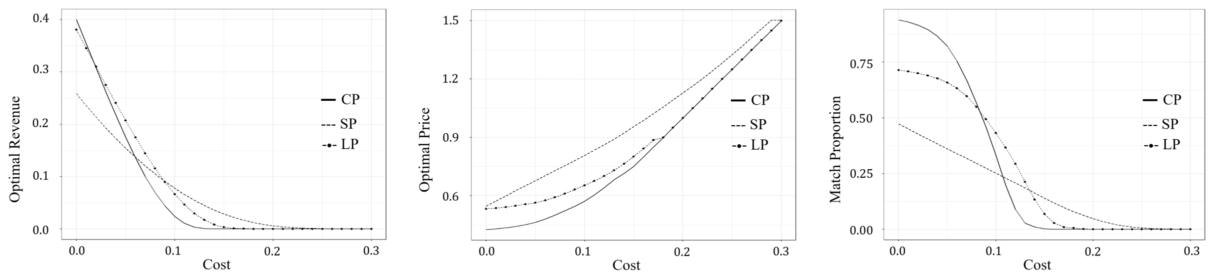

Note. Here we plot the profit, optimal prices, and match proportions under SP, CP and period L pricing with $L = T/7$, when valuations are drawn from mixed Log-normal distribution, where the two log-normal distributions Log-normal(0, 0.1) and Log-normal(0.2, 0.1), and the mixed probability is 0.6, and $T = 50$, $\delta = 0.8$, $q = 0.2$, and c varies. In the left panel, we plot the profits of SP and CP, and the note relative profit ordering switches from $\mathcal{R}_{CP} > \mathcal{R}_{SP}$ when $c \leq 0.06$, to $\mathcal{R}_{CP} > \mathcal{R}_{SP}$ for $c > 0.06$. In the middle panel, we plot the normalized optimal long period, short period and period L prices. In the right panel, we plot the proportion of the market that gets matched under SP and CP and note that $\mathcal{M}_{CP}$ dominates $\mathcal{M}_{SP}$ for $c \leq 0.105$.

Appendix C Rating System Design

C.1 Omitted Examples

In this section, we include examples omitted from the main body.

Example 9 (Theorem 12 is tight)**.** *Fix some r, a minimum level of service quality τ, and suppose the link function is $h(x) = x$, i.e., a linear function where $a = 0$ and $b = 1$. When the provider's expected quality is $q = 1$ with probability 1 (i.e. G is a point mass distribution at 1), we can plug G into Eq.* (18) *and see that the profit of any f is*

$$\max_{\substack{\beta+f/(b-a)=\tau,\\ \beta,f\in[0,1]}} r\int_{\tau}^{1} f\left(\frac{a+\beta(b-a)}{1-(q-\beta)(b-a)}\right) g(q)dq = \max_{\substack{\beta+f=\tau,\\ \beta,f\in[0,1]}} r\int_{\tau}^{1} f\left(\frac{a+\beta}{1-(q-\beta)}\right) g(q)dq = \max_{\substack{\beta+f=\tau,\\ \beta,f\in[0,1]}} rf.$$

Then, it's easy to see the optimal choice of the service fee for this G is $f^* = \tau$ yielding profit $r\tau$.

On the other hand, consider the profit generated by the choice of f and β prescribed by Theorem 12. This choice of f is $f = \frac{\left(1+2\sqrt{1-\tau}-\tau\right)\tau}{2\left(1+\sqrt{1-\tau}-\tau\right)}$ and yields profit,

$$r\int_{\tau}^{1} f\left(\frac{a+\beta(b-a)}{1-(q-\beta)(b-a)}\right) g(q)dq = rf = \frac{r\left(1+2\sqrt{1-\tau}-\tau\right)\tau}{2\left(1+\sqrt{1-\tau}-\tau\right)}$$

The corresponding ratio of the two profits is then

$$\frac{rf}{rf^*} = \frac{\left(1+2\sqrt{1-\tau}-\tau\right)}{2\left(1+\sqrt{1-\tau}-\tau\right)} = \frac{2+\sqrt{1-\tau}}{2(\sqrt{1-\tau}+1)},$$

which is the approximation ratio in Theorem 12, and converges to 3/4 as τ tends to 0.

C.2 Omitted Proofs

In this section, we include the proofs of all results omitted from the main body.

C.2.1 Omitted Proofs from Section 4.2

Proof of Theorem 10. In this proof we will show that if w_t is myopically chosen at every time step to minimize the maximum of the no-switch and switch error, then the sequence $\{w_i\}_{i=1}^{\infty}$ always converges to a moving average system. To prove this, first we will show that the myopically optimal choice of w_t always falls into one of three cases: in case 1 w_t minimizes the no-switch error, in case 2 w_t minimizes the switch error, and in case 3 w_t is such that the no-switch and switch error are equal. Having established the three cases for w_t, we will then show that eventually w_t will always stay in case 2 or case 3 i.e. there exists some t^* such that for all $t > t^*$, w_t is in case 2 (or case 3). Finally to complete the proof, we will compute the limit of $\frac{w_t}{\sum_{i=1}^{t} w_i}$ where the w_i are eventually always in case 2 (or eventually always in case 3) and show they converge to a moving average system as desired.

Expression of the MSE. To begin we derive the expression for the MSE when a switch does not or does occur. For ease of exposition, we define the following functions:

$$e_1^t(w_t) := \mathbb{E}\left[\left(\frac{\sum_{i=1}^{t} w_i s_i}{\sum_{i=1}^{t} w_i} - \mu_t\right)^2 \middle| \{s_i\}_{i=1}^{t} \sim F_1\right],$$

$$e_2^t(w_t) := \mathbb{E}\left[\left(\frac{\sum_{i=1}^{t} w_i s_i}{\sum_{i=1}^{t} w_i} - \mu_t\right)^2 \middle| \{s_i\}_{i=1}^{t-1} \sim F_1, s_t \sim F_2\right].$$

When no switch occurs at time t, all $s_i \sim F_1$ and the expected MSE of any oblivious rating system with weights $\{w_i\}_{i=1}^{t}$ is:

$$\begin{aligned}
e_1^t(w_t) &= \mathbb{E}\left[\left(\frac{\sum_{i=1}^{t} w_i s_i}{\sum_{i=1}^{t} w_i} - \mu_1\right)^2\right] \\
&= \mathbb{E}\left[\left(\frac{\sum_{i=1}^{t} w_i (s_i - \mu_1)}{\sum_{i=1}^{t} w_i}\right)^2\right] \\
&= \mathbb{E}\left[\frac{\sum_{i=1}^{t} w_i^2 (s_i - \mu_1)^2 + \sum_{i=1}^{t-1}\sum_{j=i+1}^{t} 2 w_i w_j (s_i - \mu_1)(s_j - \mu_1)}{\left(\sum_{i=1}^{t} w_i\right)^2}\right] \\
&= \frac{\sum_{i=1}^{t} w_i^2 \mathbb{E}\left[(s_i - \mu_1)^2\right] + \sum_{i=1}^{t-1}\sum_{j=i+1}^{t} 2 w_i w_j \mathbb{E}\left[(s_i - \mu_1)(s_j - \mu_1)\right]}{\left(\sum_{i=1}^{t} w_i\right)^2} \\
&= \frac{\sum_{i=1}^{t} w_i^2 \mathbb{E}\left[(s_i - \mu_1)^2\right]}{\left(\sum_{i=1}^{t} w_i\right)^2} \\
&= \frac{\sum_{i=1}^{t} w_i^2 \sigma_1^2}{\left(\sum_{i=1}^{t} w_i\right)^2} \qquad (39)
\end{aligned}$$

where the second equality follows from $\sum w_i\mu_1/\sum w_i = \mu_1$, the third equality by squaring, the fourth equality by grouping terms. The fifth equality follows from $\mathbb{E}[(s_i-\mu_1)(s_j-\mu_1)] = \mathbb{E}[(s_i-\mu_1)]\mathbb{E}[(s_j-\mu_1)] = 0$ since s_i and s_j are independent and $\mathbb{E}[s_i] = \mu_1$.

When the service quality does switch at t, $s_i \sim F_1$ for $i < t$, $s_t \sim F_2$, the expected MSE of any oblivious rating system with weights $\{w_i\}_{i=1}^t$ is:

$$
\begin{aligned}
e_2^t(w_t) &= \mathbb{E}\left[\left(\frac{\sum_{i=1}^t w_i s_i}{\sum_{i=1}^t w_i} - \mu_2\right)^2\right] \\
&= \mathbb{E}\left[\left(\frac{\sum_{i=1}^t w_i(s_i-\mu_2)}{\sum_{i=1}^t w_i}\right)^2\right] \\
&= \mathbb{E}\left[\left(\frac{\sum_{i=1}^{t-1} w_i(s_i-\mu_1) + \sum_{i=1}^{t-1} w_i(\mu_1-\mu_2) + w_t(s_t-\mu_2)}{\sum_{i=1}^t w_i}\right)^2\right],
\end{aligned} \tag{40}
$$

where the third equality follows by decomposing μ_2 into $\mu_2 = \mu_1 - (\mu_2 - \mu_1)$ and grouping terms. Now, expanding Eq. (40) into three parts we obtain:

$$
\begin{aligned}
\text{Eq. (40)} = &\ \mathbb{E}\left[\frac{\left(\sum_{i=1}^{t-1} w_i(s_i-\mu_1) + \sum_{i=1}^{t-1} w_i(\mu_1-\mu_2)\right)^2}{\left(\sum_{i=1}^t w_i\right)^2}\right] + \mathbb{E}\left[\frac{w_t^2(s_t-\mu_2)^2}{\left(\sum_{i=1}^t w_i\right)^2}\right] \\
&+ \mathbb{E}\left[\frac{2w_t(s_t-\mu_2)\left(\sum_{i=1}^{t-1} w_i(s_i-\mu_1) + \sum_{i=1}^{t-1} w_i(\mu_1-\mu_2)\right)}{\left(\sum_{i=1}^t w_i\right)^2}\right],
\end{aligned} \tag{41}
$$

Note in Eq. (41) the third term vanishes, i.e.,

$$
\mathbb{E}\left[\frac{2w_t(s_t-\mu_2)\left(\sum_{i=1}^{t-1} w_i(s_i-\mu_1) + \sum_{i=1}^{t-1} w_i(\mu_1-\mu_2)\right)}{\left(\sum_{i=1}^t w_i\right)^2}\right] = 0,
$$

since $\mathbb{E}[s_t - \mu_2] = 0$, and the s_i terms are independent with s_t. Now, for the two remaining terms from Eq. (41):

$$
\begin{aligned}
&= \mathbb{E}\left[\frac{\left(\sum_{i=1}^{t-1} w_i(s_i-\mu_1) + \sum_{i=1}^{t-1} w_i(\mu_1-\mu_2)\right)^2}{\left(\sum_{i=1}^{t} w_i\right)^2}\right] + \mathbb{E}\left[\frac{w_t^2(s_t-\mu_2)^2}{\left(\sum_{i=1}^{t} w_i\right)^2}\right] \\
&= \mathbb{E}\left[\frac{\left(\sum_{i=1}^{t-1} w_i(s_i-\mu_1)\right)^2 + \left(\sum_{i=1}^{t-1} w_i(\mu_1-\mu_2)\right)^2 + 2\left(\sum_{i=1}^{t-1} w_i(s_i-\mu_1)\right)\left(\sum_{i=1}^{t-1} w_i(\mu_1-\mu_2)\right)}{\left(\sum_{i=1}^{t} w_i\right)^2}\right] + \frac{w_t^2\sigma_2^2}{\left(\sum_{i=1}^{t} w_i\right)^2} \\
&= \mathbb{E}\left[\frac{\left(\sum_{i=1}^{t-1} w_i(s_i-\mu_1)\right)^2 + \left(\sum_{i=1}^{t-1} w_i(\mu_1-\mu_2)\right)^2}{\left(\sum_{i=1}^{t} w_i\right)^2}\right] + \frac{w_t^2\sigma_2^2}{\left(\sum_{i=1}^{t} w_i\right)^2} \\
&= \frac{\sum_{i=1}^{t-1} w_i^2\mathbb{E}\left[(s_i-\mu_1)^2\right] + \left(\sum_{i=1}^{t-1} w_i(\mu_1-\mu_2)\right)^2}{\left(\sum_{i=1}^{t} w_i\right)^2} + \frac{w_t^2\sigma_2^2}{\left(\sum_{i=1}^{t} w_i\right)^2} \\
&= \frac{\sum_{i=1}^{t-1} w_i^2\sigma_1^2 + w_t^2\sigma_2^2 + \left(\sum_{i=1}^{t-1} w_i(\mu_1-\mu_2)\right)^2}{\left(\sum_{i=1}^{t} w_i\right)^2},
\end{aligned}
\tag{42}
$$

where the second equality follows from squaring, and the third equality follows since $\mathbb{E}[s_i - \mu_1] = 0$ for $i < t$, and all the other equalities follow from simplifying. Now we are ready to consider the cases for w_t when it is chosen to minimize the max of the no switch and switch error.

Possible cases for w_t Taking the above expressions for the error when a switch does not and does occur together, and fixing the previous weights $\{w_i\}_{i=1}^{t-1}$, we aim to minimize the maximum error at each time t, i.e., find w_t such

$$
w_t = \arg\min\left\{\max\left\{\frac{\sum_{i=1}^{t} w_i^2\sigma_1^2}{\left(\sum_{i=1}^{t} w_i\right)^2}, \frac{\sum_{i=1}^{t-1} w_i^2\sigma_1^2 + w_t^2\sigma_2^2 + \left(\sum_{i=1}^{t-1} w_i(\mu_1-\mu_2)\right)^2}{\left(\sum_{i=1}^{t} w_i\right)^2}\right\}\right\},
$$

where the first term is $e_1^t(w_t)$, and the second term is $e_2^t(w_t)$. However, where the maximum error occurs depends on the interactions between $e_1^t(w_t)$ and $e_2^t(w_t)$. To understand where the min-max solution can occur, we first will study the first order derivatives of $e_1^t(w_t)$ and $e_2^t(w_t)$ with respect to w_t. Specifically, the first order derivative of $e_1^t(w_t)$ is

$$
\begin{aligned}
\frac{\partial e_1^t(w_t)}{\partial w_t} &= \frac{2w_t\sigma_1^2\left(\sum_{i=1}^t w_i\right)^2 - 2\left(\sum_{i=1}^t w_i\right)\left(\sum_{i=1}^t w_i^2\sigma_1^2\right)}{\left(\sum_{i=1}^t w_i\right)^4} \\
&= \frac{2w_t\sigma_1^2\left(\sum_{i=1}^t w_i\right) - 2\left(\sum_{i=1}^t w_i^2\sigma_1^2\right)}{\left(\sum_{i=1}^t w_i\right)^3} \\
&= \frac{2w_t\sigma_1^2\left(\sum_{i=1}^{t-1} w_i\right) - 2\left(\sum_{i=1}^{t-1} w_i^2\sigma_1^2\right)}{\left(\sum_{i=1}^t w_i\right)^3},
\end{aligned}
$$

where the second equality follows from canceling $\sum_{i=1}^t w_i$ and the third equality follows from canceling $w_t^2\sigma_1^2$. Moreover, we observe the numerator of $\frac{\partial e_1^t(w_t)}{\partial w_t}$ is linear in w_t. When $w_t = 0$, we have

$$
\frac{\partial e_1^t(w_t)}{\partial w_t} = -\frac{2\left(\sum_{i=1}^{t-1} w_i^2\sigma_1^2\right)}{\left(\sum_{i=1}^t w_i\right)^3} < 0,
$$

when w_t tends to infinity, we have

$$
\lim_{t\to\infty} \frac{\partial e_1^t(w_t)}{\partial w_t} = \lim_{t\to\infty} \frac{2w_t\sigma_1^2\left(\sum_{i=1}^{t-1} w_i\right) - 2\left(\sum_{i=1}^{t-1} w_i^2\sigma_1^2\right)}{\left(\sum_{i=1}^t w_i\right)^3} > 0.
$$

Therefore, the no switching error $e_1^t(w_t)$ is first decreasing in w_t and then increasing in w_t, there is a unique minimum of $e_1^t(w_t)$.

Similarly, we consider the first order derivative of $e_2^t(w_t)$ with respect to w_t, which is

$$
\begin{aligned}
\frac{\partial e_2^t(w_t)}{\partial w_t} &= \frac{2w_t\sigma_2^2\left(\sum_{i=1}^t w_i\right)^2 - 2\left(\sum_{i=1}^t w_i\right)\left(\sum_{i=1}^{t-1} w_i^2\sigma_1^2 + \left(\sum_{i=1}^{t-1} w_i(\mu_1-\mu_2)\right)^2 + w_t^2\sigma_2^2\right)}{\left(\sum_{i=1}^t w_i\right)^4} \\
&= \frac{2w_t\sigma_2^2\left(\sum_{i=1}^t w_i\right) - 2\left(\sum_{i=1}^{t-1} w_i^2\sigma_1^2 + \left(\sum_{i=1}^{t-1} w_i(\mu_1-\mu_2)\right)^2 + w_t^2\sigma_2^2\right)}{\left(\sum_{i=1}^t w_i\right)^3} \\
&= \frac{2w_t\sigma_2^2\left(\sum_{i=1}^{t-1} w_i\right) + 2w_t^2\sigma_2^2 - 2\left(\sum_{i=1}^{t-1} w_i^2\sigma_1^2 + \left(\sum_{i=1}^{t-1} w_i(\mu_1-\mu_2)\right)^2 + w_t^2\sigma_2^2\right)}{\left(\sum_{i=1}^t w_i\right)^3} \\
&= \frac{2w_t\sigma_2^2\left(\sum_{i=1}^{t-1} w_i\right) - 2\left(\sum_{i=1}^{t-1} w_i^2\sigma_1^2 + \left(\sum_{i=1}^{t-1} w_i(\mu_1-\mu_2)\right)^2\right)}{\left(\sum_{i=1}^t w_i\right)^3},
\end{aligned}
$$

where the first equality follows from taking the derivative, the second equality follows from canceling $\sum_{i=1}^t w_i$, and the last two equality follows from rearrangement and simplifying.

Again, we see from $\frac{\partial e_2^t(w_t)}{\partial w_t}$ is linear w_t, the switch error $e_1^t(w_t)$ is also first decreasing in w_t and then increasing in w_t, there is a unique minimum of $e_2^t(w_t)$.

From the first order derivatives, we see both $e_1^t(w_t)$ and $e_2^t(w_t)$ are first decreasing and then increasing in w_t. Thus, there are three possible cases for where the choice w_t such that it minimizes the maximum MSE may occur (for a depiction of these cases see Fig. 27).

Case 1. The min-max MSE occurs when there is no quality switch (cf. the left panel of Fig. 27), i.e.,

$$w_t^{(1)} := \arg\min_{w_t} \frac{\sum_{i=1}^{t} w_i^2 \sigma_1^2}{\left(\sum_{i=1}^{t} w_i\right)^2} = \left(\sum_{i=1}^{t-1} \left(\frac{w_i}{\sum_{j=1}^{t-1} w_j}\right)^2\right) \sum_{i=1}^{t-1} w_i, \tag{43}$$

where the second equality follows from solving the first order condition $\frac{\partial e_1^t(w_t)}{\partial w_t} = 0$.

Case 2. The min-max MSE occurs when there is a service quality switch (cf. the middle panel of Fig. 27), i.e.,

$$\begin{aligned} w_t^{(2)} &:= \arg\min_{w_t} \frac{\sum_{i=1}^{t-1} w_i^2 \sigma_1^2 + w_t^2 \sigma_2^2 + \left(\sum_{i=1}^{t-1} w_i(\mu_1 - \mu_2)\right)^2}{\left(\sum_{i=1}^{t} w_i\right)^2} \\ &= \left(\frac{(\mu_1 - \mu_2)^2}{\sigma_2^2} + \frac{\sigma_1^2}{\sigma_2^2} \sum_{i=1}^{t-1} \left(\frac{w_i}{\sum_{i=1}^{t-1} w_i}\right)^2\right) \sum_{i=1}^{t-1} w_i, \end{aligned} \tag{44}$$

where the second equality follows from solving the first order condition $\frac{\partial e_2^t(w_t)}{\partial w_t} = 0$.

Case 3. The min-max MSE occurs when the no-switching error equals to switching error (cf. the right panel of Fig. 27), i.e., $w_t^{(3)}$ is the solution of the following equation,

$$\frac{\sum_{i=1}^{t} w_i^2 \sigma_1^2}{\left(\sum_{i=1}^{t} w_i\right)^2} = \frac{\sum_{i=1}^{t-1} w_i^2 \sigma_1^2 + w_t^2 \sigma_2^2 + \left(\sum_{i=1}^{t-1} w_i(\mu_1 - \mu_2)\right)^2}{\left(\sum_{i=1}^{t} w_i\right)^2}.$$

Solving the equation above, we obtain

$$w_t^{(3)} := \frac{|\mu_1 - \mu_2|}{\sqrt{\sigma_1^2 - \sigma_2^2}} \sum_{i=1}^{t-1} w_i, \tag{45}$$

when $\sigma_1 > \sigma_2$, otherwise, there is no solution.

Now we will argue that the min-max MSE can only occur either in case 2 or case 3 in the long run. To see this, suppose the maximum MSE is incurred from no switching (i.e. in case 1). Then the weights of the myopically optimal oblivious rating system must be chosen to always minimize $\frac{\sum_{i=1}^{t} w_i^2 \sigma_1^2}{\left(\sum_{i=1}^{t} w_i\right)^2}$ in this case. However, it is not hard to see that a simple moving average rating system will not only minimize this term, but asymptotically make it vanish since the no-switch MSE of the simple average rating system $\frac{\sum_{i=1}^{t} w_i^2 \sigma_1^2}{\left(\sum_{i=1}^{t} w_i\right)^2}$ converges to 0 when t goes to infinity.

However, the assumption that $\mu_1 \neq \mu_2$ ensures that the switching error can never be made 0. Namely, the switching error at any time t is always lower bounded by $\min\left\{\frac{\sigma_2^2}{4}, \frac{(\mu_1-\mu_2)^2}{4}\right\} > 0$ since

$$
\begin{aligned}
\frac{\sum_{i=1}^{t-1} w_i^2 \sigma_1^2 + w_t^2 \sigma_2^2 + \left(\sum_{i=1}^{t-1} w_i(\mu_1 - \mu_2)\right)^2}{\left(\sum_{i=1}^{t} w_i\right)^2} &\geq \frac{w_t^2 \sigma_2^2 + \left(\sum_{i=1}^{t-1} w_i(\mu_1 - \mu_2)\right)^2}{\left(\sum_{i=1}^{t} w_i\right)^2} \\
&= \left(\frac{w_t}{\sum_{i=1}^{t} w_i}\right)^2 \sigma_2^2 + \left(1 - \frac{w_t}{\sum_{i=1}^{t} w_i}\right)^2 (\mu_1 - \mu_2)^2 \\
&\geq \min\left\{\frac{\sigma_2^2}{4}, \frac{(\mu_1 - \mu_2)^2}{4}\right\},
\end{aligned}
$$

where the first inequality follows from $w_i^2 \sigma_1^2 \geq 0$ for $i < t$, and the second inequality follows since when $\frac{w_t}{\sum_{i=1}^{t} w_i} \geq \frac{1}{2}$, then $\left(\frac{w_t}{\sum_{i=1}^{t} w_i}\right)^2 \sigma_2^2 \geq \frac{\sigma_1^2}{4}$, when $\frac{w_t}{\sum_{i=1}^{t} w_i} \leq \frac{1}{2}$, then $\left(1 - \frac{w_t}{\sum_{i=1}^{t} w_i}\right)^2 (\mu_1 - \mu_2)^2 \geq \frac{(\mu_1-\mu_2)^2}{4}$.

Thus, there are two cases for minimizing the maximum MSE in the long run, either w_t will minimize the switching error or w_t will balance the MSE between when a service quality switch does or does not occur. Now are ready to characterize the long run behavior of w_t when it is always chosen to minimize the maximum error. We will show that eventually w_t will always be in case 2 or always be in case 3 (and thus no oscillates between the cases). To prove, we will need to analyze several sub-cases based on the specification of the parameters σ_1, σ_2, μ_1, and μ_2.

Figure 27: Cases for the min-max choice of w_t.

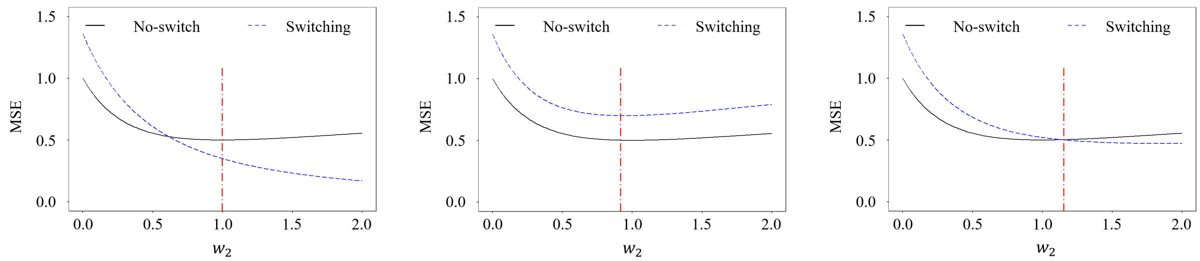

Note. Here we plot the switching error and non-switching error for an instance where $w_1 = 1$ and $w_2 \in [0, 2]$, $\mu_1 = 0.7$, $\mu_2 = 0.1$ and $\sigma_1 = 1$, and where σ_2 varies across panels such that $\sigma_2 = 0.2, 1.2, 0.85$ in the left, middle, and right panels, respectively. In all panels the choice of w_2 that achieves the min-max MSE is denoted by the red line. In the left panel, we note the min-max error occurs at value that minimizes the no-switching error (case 1). In the middle panel, we note the min-max error occurs at the value that minimizes the switching error (case 2). In the right panel, we note the min-max error occurs at value that equates the no-switching and switching error (case 3).

When $\sigma_1 \leq \sigma_2$: Suppose $\sigma_1 \leq \sigma_2$, when this holds the switching error always dominates the no-switching error i.e.,

$$\frac{\sum_{i=1}^{t-1} w_i^2 \sigma_1^2 + w_t^2 \sigma_2^2 + \left(\sum_{i=1}^{t-1} w_i (\mu_1 - \mu_2)\right)^2}{\left(\sum_{i=1}^{t} w_i\right)^2} \geq \frac{\sum_{i=1}^{t} w_i^2 \sigma_1^2}{\left(\sum_{i=1}^{t} w_i\right)^2}$$

where the inequality follows from $\sigma_2^2 \geq \sigma_1^2$ and $\left(\sum_{i=1}^{t-1} w_i (\mu_1 - \mu_2)\right)^2 \geq 0$. Thus, when $\sigma_2 \geq \sigma_1$, the min-max MSE can only occur for w_t in case 2. Now we want to show the convergence of $\lim_{t\to\infty} \frac{w_t^{(2)}}{\sum_{i=1}^{t} w_i^{(2)}}$.

Let $\eta_t := \sum_{i=1}^{t} \frac{w_i^2}{\left(\sum_{i=1}^{t} w_i\right)^2}$, and let η^* be the solution of the following equation,

$$\frac{(\mu_1 - \mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right) \eta^* - \frac{2\eta^*}{1 - \eta^*} = 0 \tag{46}$$

for $\eta^* \in (0, 1)$. It can be verified that $\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right) \eta - \frac{2\eta}{1-\eta}$ is concave in η for $\eta \in (0, 1)$, and when $\eta = 0$, $\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right) \eta - \frac{2\eta}{1-\eta} = \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} > 0$, when η tends to 1, $\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right) \eta - \frac{2\eta}{1-\eta}$

tends to negative infinity. Therefore, the solution of Eq. (46) is unique, and η^* is well defined. Further, when $\eta \le \eta^*$, $\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta - \frac{2\eta}{1-\eta} \ge 0$; when $\eta \ge \eta^*$, $\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta - \frac{2\eta}{1-\eta} \le 0$. In the following, we will show that for w_t defined in the recursion of Eq. (44), $|\eta_t - \eta^*|$ is decreasing.

Suppose at time t we have $\eta_t > \eta^*$. By the recursion in Eq. (44), the weight w_{t+1} at time $t+1$ is

$$w_{t+1} = \left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)\sum_{i=1}^{t} w_i.$$

Then, η_{t+1} will be

$$\begin{aligned}
\eta_{t+1} &= \sum_{i=1}^{t+1} \frac{w_i^2}{\left(\sum_{j=1}^{t+1} w_j\right)^2} \\
&= \frac{\sum_{i=1}^{t} w_i^2 + \left(\left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)\sum_{i=1}^{t} w_i\right)^2}{\left(\sum_{j=1}^{t} w_j + \left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)\sum_{i=1}^{t} w_i\right)^2} \\
&= \frac{\sum_{i=1}^{t} w_i^2 / \left(\sum_{j=1}^{t} w_j\right)^2 + \left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2}{\left(1 + \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2} \\
&= \frac{\eta_t + \left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2}{\left(1 + \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2},
\end{aligned}$$

where the first equality follows from the definition of η_t, the second equality follows plugging in $w_{t+1} = \left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)\sum_{i=1}^{t} w_i$, the third equality follows from rearranging, and the fourth equality follows from the definition of η_t. Next, we show if $\eta_t \ge \eta^*$, then $\eta_t \ge \eta_{t+1} \ge \eta^*$. To show $\eta_t \ge \eta_{t+1}$, we need to show

$$\eta_t \ge \frac{\eta_t + \left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2}{\left(1 + \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2},$$

which is equivalent to

$$\eta_t + \left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2 \le \eta_t\left(1 + \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2.$$

Note that the difference between $\eta_t + \left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2$ and $\eta_t\left(1 + \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2$ is

$$
\begin{aligned}
&\eta_t + \left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2 - \eta_t\left(1 + \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2 \\
&= \left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2(1-\eta_t) - 2\eta_t\left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right) \\
&= (1-\eta_t)\left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)\left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t - \frac{2\eta_t}{1-\eta_t}\right) \\
&\leq 0,
\end{aligned}
$$

where the first equality follows from canceling η_t, the second equality follows from taking $(1-\eta_t)\left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)$ out, the inequality follows from when $\eta_t \in [\eta^*, 1)$, we have $\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t - \frac{2\eta_t}{1-\eta_t} \leq 0$. Thus, when $\eta_t \geq \eta^*$, then $\eta_{t+1} \leq \eta_t$.

Next, we show if $\eta_t \geq \eta^*$, then $\eta_{t+1} \geq \eta^*$, which is equivalent to show

$$
\frac{\eta_t + \left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2}{\left(1 + \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2} \geq \eta^*.
$$

Specifically, we only need to show

$$
\eta_t + \left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2 \geq \eta^*\left(1 + \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2.
$$

Note that the difference between $\eta_t + \left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2$ and $\eta^*\left(1 + \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2$ is

$$
\begin{aligned}
&\eta_t + \left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2 - \eta^*\left(1 + \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2 \\
&= (\eta_t - \eta^*) + \left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)^2(1-\eta^*) - 2\eta^*\left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right) \\
&= (\eta_t - \eta^*) + (1-\eta^*)\left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)\left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t - \frac{2\eta^*}{1-\eta^*}\right) \\
&\geq (1-\eta^*)\left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)\left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta^* - \frac{2\eta^*}{1-\eta^*}\right) \\
&\geq 0,
\end{aligned}
$$

where the first equality follows from grouping terms, the second equality follows from taking $(1-\eta^*)\left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2}+\left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_t\right)$ out, the first inequality follows from $\eta_t \geq \eta^*$, and the second inequality follows from definition of η^*, i.e., $\frac{(\mu_1-\mu_2)^2}{\sigma_2^2}+\left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta^* - \frac{2\eta^*}{1-\eta^*} = 0$.

Therefore, combining the analysis above, we conclude if $\eta_t \geq \eta^*$, then $0 \leq \eta_{t+1}-\eta^* \leq \eta_t - \eta^*$. Similarly, follow the same structure, it can be shown that if $\eta_t \leq \eta^*$, then $\eta_t \leq \eta_{t+1} \leq \eta^*$. Thus $|\eta_t - \eta^*|$ is decreasing in t, and the decreasing is strict if $\eta_t \neq \eta^*$. By fixed point theorem, η_t converges to a constant. Thus, there exists a constant α such that

$$\lim_{t\to\infty} \frac{w_t^{(2)}}{\sum_{i=1}^t w_i^{(2)}} = \alpha,$$

as desired.

When $\sigma_1 > \sigma_2$: Now consider the situation when $\sigma_1 > \sigma_2$. This case is more involved than the previous case. As we will see, both case 2 and case 3 solutions for w_t are possible and will break down into sub-cases w_t is eventually always in case 2 (where the limit computed above implies that the sequence converges to a moving average), or in case 3 where we will need to check a new limit.

Now, recall the definition of $e_1^t(w_t)$ and $e_2^t(w_t)$. When $\sigma_1 > \sigma_2$, we can see that

$$\begin{aligned}
e_1^t(0) &= \frac{\sum_{i=1}^{t-1} w_i^2}{\left(\sum_{i=1}^{t-1} w_i\right)^2}, \\
e_2^t(0) &= \frac{\sum_{i=1}^{t-1} w_i^2}{\left(\sum_{i=1}^{t-1} w_i\right)^2} + (\mu_1-\mu_2)^2 > e_1^t(0),
\end{aligned}$$

where the inequality follows from the assumption that $\mu_1 \neq \mu_2$. Since both $e_1^t(w_t)$ and $e_2^t(w_t)$ are first decreasing and then increasing in w_t, in this scenario where $\sigma_1 > \sigma_2$ the conditions for the three cases can be simplified as:

1. The min-max occurs in case 1 if

$$w_t^{(3)} \leq w_t^{(1)} \leq w_t^{(2)}.$$

2. The min-max occurs in case 2 if

$$w_t^{(1)} \leq w_t^{(2)} \leq w_t^{(3)}.$$

3. The min-max occurs in case 3 if

$$w_t^{(1)} \leq w_t^{(3)} \leq w_t^{(2)}.$$

based on the crossing of e_1 and e_2, and where recall the superscripts denote that choice of w_t specific to case 1, 2 and 3, respectively. Note $w_t^{(2)}$ is always larger than $w_t^{(1)}$ by definition. By our previous argument, we know the case of w_t it will not converge to case 1 in the long run. Thus we only need to consider case 2 and case 3. For simplicity of notations, let $\kappa = |\mu_1 - \mu_2|/\sqrt{\sigma_1^2 - \sigma_2^2}$, in the following we will show that when

$$\kappa \leq \frac{(\mu_1 - \mu_2)^2}{\sigma_2^2} + \left(\frac{\kappa}{2+\kappa}\right)\frac{\sigma_1^2}{\sigma_2^2}$$

the case of w_t converges to case 3, and otherwise the case of w_t converges to case 2, see Fig. 28 for an depiction.

Figure 28: Convergence of Cases.

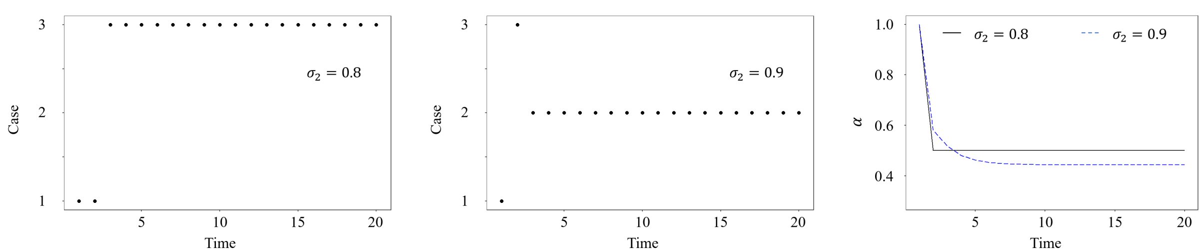

Note. Here we plot the convergence of cases over time for an example instance where $\mu_1 = 0.7$, $\mu_2 = 0.1$, $\sigma_1 = 1$, and where σ_2 varies across panels such that $\sigma_2 = 0.8, 0.9$ in the left and middle panels respectively. In the left panel, we plot the optimal cases of w_t as t goes from 1 to 20 when $\sigma_2 = 0.8$. Here $\kappa = 1 \leq \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\kappa}{2+\kappa}\right)\frac{\sigma_1^2}{\sigma_2^2} \approx 1.083$, and we can see the optimal case converges to case 3. In the middle panel, we plot the optimal cases of w_t as t goes from 1 to 20 when $\sigma_2 = 0.9$. Here $\kappa \approx 1.376 \geq \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\kappa}{2+\kappa}\right)\frac{\sigma_1^2}{\sigma_2^2} \approx 0.948$, and we can see the optimal case converges to case 2. In the right panel, we plot the convergence of the ratio of the most recent weight to the previous weights, $\frac{w_t}{\sum_{i=1}^{t} w_i}$, as t goes from 1 to 20 for the two panels.

When $\sigma_1 > \sigma_2$ and $\kappa \leq \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\kappa}{2+\kappa}\right)\frac{\sigma_1^2}{\sigma_2^2}$: Here show that when $\sigma_1 > \sigma_2$ and $\kappa \leq \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\kappa}{2+\kappa}\right)\frac{\sigma_1^2}{\sigma_2^2}$, the case of w_t converges to case 3. Note that the condition for the solution to be in case 3 is

$$w_t^{(1)} \leq w_t^{(3)} \leq w_t^{(2)},$$

which means at time t

$$\kappa = \frac{|\mu_1 - \mu_2|}{\sqrt{\sigma_1^2 - \sigma_2^2}} \leq \frac{(\mu_1 - \mu_2)^2}{\sigma_2^2} + \frac{\sigma_1^2}{\sigma_2^2}\left(\sum_{i=1}^{t-1}\left(\frac{w_i}{\sum_{j=1}^{t-1} w_j}\right)^2\right), \tag{47}$$

where the condition follows directly from plugging in the values for $w_t^{(2)}, w_t^{(3)}$, and canceling $\sum_{i=1}^{t-1} w_i$ across. With the condition in Eq. (47), we will show that plugging in $w_t = w_t^{(3)}$, then in next step $w_{t+1}^{(3)} \leq w_{t+1}^{(2)}$ which implies case 3 is absorbing. Specifically, since the solution of w_t is in case 3,

$$w_t = w_t^{(3)} = \kappa \sum_{i=1}^{t-1} w_i.$$

Thus, $w_{t+1}^{(3)}$ is

$$w_{t+1}^{(3)} = \kappa \sum_{i=1}^{t} w_i = \kappa(1+\kappa)\sum_{i=1}^{t-1} w_i.$$

Similarly, $w_{t+1}^{(2)}$ is

$$\begin{aligned}
w_{t+1}^{(2)} &= \left(\frac{(\mu_1 - \mu_2)^2}{\sigma_2^2} + \frac{\sigma_1^2}{\sigma_2^2}\left(\sum_{i=1}^{t}\left(\frac{w_i}{\sum_{j=1}^{t} w_j}\right)^2\right)\right)\sum_{i=1}^{t} w_i \\
&= \left(\frac{(\mu_1 - \mu_2)^2}{\sigma_2^2} + \frac{\sigma_1^2}{\sigma_2^2}\left(\sum_{i=1}^{t}\left(\frac{w_i}{(1+\kappa)\sum_{j=1}^{t-1} w_j}\right)^2\right)\right)(1+\kappa)\sum_{i=1}^{t-1} w_i \\
&= \left(\frac{(\mu_1 - \mu_2)^2}{\sigma_2^2} + \frac{\sigma_1^2}{\sigma_2^2}\left(\sum_{i=1}^{t-1}\left(\frac{w_i}{(1+\kappa)\sum_{j=1}^{t-1} w_j}\right)^2 + \frac{\kappa^2}{(1+\kappa)^2}\right)\right)(1+\kappa)\sum_{i=1}^{t-1} w_i \\
&= \left(\frac{(\mu_1 - \mu_2)^2}{\sigma_2^2} + \frac{\sigma_1^2}{\sigma_2^2}\left(\sum_{i=1}^{t-1}\left(\frac{w_i}{(1+\kappa)\sum_{j=1}^{t-1} w_j}\right)^2 + \frac{\kappa^2}{(1+\kappa)^2}\right)\right)(1+\kappa)\sum_{i=1}^{t-1} w_i,
\end{aligned}$$

where the second equality follows from plugging $w_t = \kappa \sum_{i=1}^{t} w_i$, and all the other equalities follow from rearranging and simplifying. Now to show $w_{t+1}^{(3)} \leq w_{t+1}^{(2)}$ it is equivalent to show that

$$\kappa \leq \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \frac{\sigma_1^2}{\sigma_2^2}\left(\sum_{i=1}^{t-1}\left(\frac{w_i}{(1+\kappa)\sum_{j=1}^{t-1} w_j}\right)^2 + \frac{\kappa^2}{(1+\kappa)^2}\right).$$

Note that the difference of the right hand side minus the left, when multiplied by $(1+\kappa)^2$, is

$$\begin{aligned}
&(1+\kappa)^2\left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2}\right) + \frac{\sigma_1^2}{\sigma_2^2}\left(\sum_{i=1}^{t-1}\left(\frac{w_i}{\sum_{j=1}^{t-1} w_j}\right)^2 + \kappa^2\right) - \kappa(1+\kappa)^2\\
&= \left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2}\right) + \frac{\sigma_1^2}{\sigma_2^2}\left(\sum_{i=1}^{t-1}\left(\frac{w_i}{\sum_{j=1}^{t-1} w_j}\right)^2 + \kappa^2\right) - \kappa + \left((2\kappa+\kappa^2)\left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} - \kappa\right)\right)\\
&= \left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2}\right) + \frac{\sigma_1^2}{\sigma_2^2}\left(\sum_{i=1}^{t-1}\left(\frac{w_i}{\sum_{j=1}^{t-1} w_j}\right)^2\right) - \kappa + \left((2\kappa+\kappa^2)\left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \frac{\sigma_1^2}{\sigma_2^2}\left(\frac{\kappa^2}{2\kappa+\kappa^2}\right) - \kappa\right)\right)\\
&= \left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2}\right) + \frac{\sigma_1^2}{\sigma_2^2}\left(\sum_{i=1}^{t-1}\left(\frac{w_i}{\sum_{j=1}^{t-1} w_j}\right)^2\right) - \kappa + \left((2\kappa+\kappa^2)\left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \frac{\sigma_1^2}{\sigma_2^2}\left(\frac{\kappa}{2+\kappa}\right) - \kappa\right)\right)\\
&\geq 0
\end{aligned}$$

where the first equality follows from $(1+\kappa)^2 = 1+2\kappa+\kappa^2$, all the other equalities follow from rearranging, and the inequality follows from $\left(\frac{(\mu_1-\mu_2)^2}{\sigma_2^2}\right) + \frac{\sigma_1^2}{\sigma_2^2}\left(\sum_{i=1}^{t-1}\left(\frac{w_i}{\sum_{j=1}^{t-1} w_j}\right)^2\right) - \kappa \geq 0$ since at time t the optimal solution is in case 3, and $\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \frac{\sigma_1^2}{\sigma_2^2}\left(\frac{\kappa}{2+\kappa}\right) - \kappa \geq 0$ by assumption. Therefore, when $\kappa \leq \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \frac{\sigma_1^2}{\sigma_2^2}\left(\frac{\kappa}{2+\kappa}\right)$, case 3 will be an absorbing state: if the optimal choice of w_t is ever in case 3, it will never switch to another case. This leaves two options, either eventually w_t is always in case 2, in which case we know it converges to a moving average by above, or it is eventually always in case 3. In case 3, the ratio between $w_t^{(3)}$ and $\sum_{i=1}^{t-1} w_i^{(3)}$ is always a constant since $w_t^{(3)} \propto \sum_{i=1}^{t-1} w_i^{(3)}$. Therefore, there exists a constant α such that

$$\lim_{t\to\infty} \frac{w_t^{(3)}}{\sum_{i=1}^{t} w_i^{(3)}} = \alpha.$$

When $\sigma_1 > \sigma_2$ and $\kappa > \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\kappa}{2+\kappa}\right)\frac{\sigma_1^2}{\sigma_2^2}$: Finally consider when $\sigma_1 > \sigma_2$ and $\kappa > \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\kappa}{2+\kappa}\right)\frac{\sigma_1^2}{\sigma_2^2}$. In this scenario we will show that if the case of w_t is ever case 2, then it will stay in case 2. To see this, suppose w_t is in case 2, or equivalently that the following holds

$$\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \frac{\sigma_1^2}{\sigma_2^2}\left(\sum_{i=1}^{t-1}\left(\frac{w_i}{\sum_{j=1}^{t-1} w_j}\right)^2\right) \leq \frac{|\mu_1-\mu_2|}{\sqrt{\sigma_1^2-\sigma_2^2}},$$

then, we will show $w_{t+1}^{(2)} \leq w_{t+1}^{(3)}$, which implies case 2 is absorbing. Recall the definition of η^*, and the contraction mapping nature of η_{t+1} when w_t is in case 2. Let κ^* be the solution of the following equation,

$$\kappa^* = \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\kappa^*}{2+\kappa^*}\right)\frac{\sigma_1^2}{\sigma_2^2}.$$

Recall that η^* is the solution of

$$\frac{2\eta^*}{1-\eta^*} = \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta^*.$$

Therefore, the relation between η^* and κ^* is

$$\kappa^* = \frac{2\eta^*}{1-\eta^*}, \quad \frac{\kappa^*}{2+\kappa^*} = \frac{\frac{2\eta^*}{1-\eta^*}}{2+\frac{2\eta^*}{1-\eta^*}} = \eta^*.$$

The two equations for κ^* and η^* are in the same form. By the concavity of $\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\kappa}{2+\kappa}\right)\frac{\sigma_1^2}{\sigma_2^2}$, if and only if $\kappa \geq \kappa^*$, the inequality $\kappa > \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\kappa}{2+\kappa}\right)\frac{\sigma_1^2}{\sigma_2^2}$ holds. Then, we have $\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta^* < \kappa$ when $\kappa > \frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\kappa}{2+\kappa}\right)\frac{\sigma_1^2}{\sigma_2^2}$.

Now, if $\eta_{t-1} = \left(\frac{w_i}{\sum_{j=1}^{t} w_j}\right)^2 \leq \eta^*$, from our previous analysis of η_t when w_t is in case 2, we know η_t will converge to η^* from below, and will never cross η^*. Therefore, at time $t+1$, we also have

$$\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \frac{\sigma_1^2}{\sigma_2^2}\left(\sum_{i=1}^{t}\left(\frac{w_i}{\sum_{j=1}^{t} w_j}\right)^2\right) \leq \frac{|\mu_1-\mu_2|}{\sqrt{\sigma_1^2-\sigma_2^2}}.$$

Similarly, if $\eta_{t-1} \geq \eta^*$, and w_t is in case 2 at time t, i.e., $\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \left(\frac{\sigma_1^2}{\sigma_2^2}\right)\eta_{t-1} < \kappa$, then η_t will converge to η^* from above, i.e., $\eta_t \leq \eta_{t-1}$. Thus, at time $t+1$, we have

$$\frac{(\mu_1-\mu_2)^2}{\sigma_2^2} + \frac{\sigma_1^2}{\sigma_2^2}\left(\sum_{i=1}^{t}\left(\frac{w_i}{\sum_{j=1}^{t} w_j}\right)^2\right) \leq \frac{|\mu_1-\mu_2|}{\sqrt{\sigma_1^2-\sigma_2^2}}.$$

Therefore, we conclude that case 2 is absorbing. The convergence result is the same as in the scenario when $\sigma_1 \leq \sigma_2$.

Combining all the cases, we conclude that $\frac{w_t}{\sum_{i=1}^{t} w_i}$ converges to a constant in the long run. This is implies that the myopically optimal oblivious rating system converges to a α-moving average rating system, since

$$\begin{aligned}
R_t(\{s_i\}_{i=1}^{t}) &= \frac{\sum_{i=1}^{t} w_i s_i}{\sum_{i=1}^{t} w_i} \\
&= \frac{\sum_{i=1}^{t-1} w_i s_i}{\sum_{i=1}^{t} w_i} + \left(\frac{w_t}{\sum_{i=1}^{t} w_i}\right) s_t \\
&= \frac{\sum_{i=1}^{t-1} w_i s_i / \sum_{i=1}^{t-1} w_i}{\sum_{i=1}^{t} w_i / \sum_{i=1}^{t-1} w_i} + \left(\frac{w_t}{\sum_{i=1}^{t} w_i}\right) s_t \\
&= R_{t-1}(\{s_i\}_{i=1}^{t-1})\left(\sum_{i=1}^{t-1} w_i / \sum_{i=1}^{t} w_i\right) + \left(\frac{w_t}{\sum_{i=1}^{t} w_i}\right) s_t \\
&= \left(1 - \frac{w_t}{\sum_{i=1}^{t} w_i}\right) R_{t-1}(\{s_i\}_{i=1}^{t-1}) + \left(\frac{w_t}{\sum_{i=1}^{t} w_i}\right) s_t \\
&\approx (1-\alpha) R_{t-1}(\{s_i\}_{i=1}^{t-1}) + \alpha s_t,
\end{aligned}$$

where the first equality is the definition of a rating system, third equality follows from dividing $\sum_{i=1}^{t-1} w_i$ for both the numerator and denominator in the first term, and the final approximation follows by the limit. □

Proof of Corollary 1. Fix some α and note that under the α-moving average rating system, the MSE in limit when no switch occurs is:

$$\begin{aligned}
\lim_{t\to\infty} e_1^t(w_t) &= \mathbb{E}\left[\left(\sum_{i=0}^{\infty}\alpha(1-\alpha)^i s_{t-i} - \mu_1\right)^2\right] \\
&= \mathbb{E}\left[\left(\sum_{i=0}^{\infty}\alpha(1-\alpha)^i (s_{t-i} - \mu_1)\right)^2\right] \\
&= \mathbb{E}\left[\sum_{i=1}^{\infty}\sum_{j=1}^{\infty}\left(\alpha(1-\alpha)^i(s_{t-i}-\mu_1)\right)\left(\alpha(1-\alpha)^j(s_{t-j}-\mu_1)\right)\right] \\
&= \sum_{i=1}^{\infty}\alpha^2(1-\alpha)^{2i}\mathbb{E}\left[(s_{t-i}-\mu_1)^2\right] \\
&= \sum_{i=1}^{\infty}\alpha^2(1-\alpha)^{2i}\sigma_1^2 \\
&= \left(\frac{\alpha}{2-\alpha}\right)\sigma_1^2,
\end{aligned}$$

where the fourth equality follows from the independence of reviews s_i and s_j for all $i \neq j$. When the service quality switch does occur in limit, the MSE for the α-moving average rating system is:

$$\begin{aligned}
\lim_{t\to\infty} e_2^t(w_t) &= \mathbb{E}\left[\left(\sum_{i=0}^{\infty}\alpha(1-\alpha)^i s_{t-i} - \mu_2\right)^2\right] \\
&= \mathbb{E}\left[\left(\sum_{i=1}^{\infty}\alpha(1-\alpha)^i (s_{t-i} - \mu_1) + (1-\alpha)(\mu_1-\mu_2) + \alpha(s_t-\mu_2)\right)^2\right] \\
&= \sum_{i=1}^{\infty}\alpha^2(1-\alpha)^{2i}\mathbb{E}[(s_{t-i}-\mu_1)^2] + (1-\alpha)^2(\mu_1-\mu_2)^2 + \alpha^2\mathbb{E}[(s_t-\mu_2)^2] \\
&= \left(\frac{\alpha}{2-\alpha}\right)\sigma_1^2 + (1-\alpha)^2(\mu_1-\mu_2)^2 + \alpha^2(\sigma_2^2-\sigma_1^2),
\end{aligned}$$

where the third equality follows the independence of s_{t-i} and s_{t-j} for all $i \neq j$.

To optimize the α-moving average rating system, we aim to calculate α such that in limit the maximum MSE when a no-switch or a switch occurs is minimized, i.e.,

$$\min_{\alpha}\max\left\{\left(\frac{\alpha}{2-\alpha}\right)\sigma_1^2, \left(\frac{\alpha}{2-\alpha}\right)\sigma_1^2 + (1-\alpha)^2(\mu_1-\mu_2)^2 + \alpha^2(\sigma_2^2-\sigma_1^2)\right\}. \tag{48}$$

Note that $\frac{\alpha}{2-\alpha}$ is always increasing in α for $\alpha \in [0,1]$, and the derivative of the second term is

$$\frac{\partial}{\partial \alpha}\left(\left(\frac{\alpha}{2-\alpha}\right)\sigma_1^2 + (1-\alpha)^2(\mu_1-\mu_2)^2 + \alpha^2(\sigma_2^2-\sigma_1^2)\right) = 2\left(\frac{\sigma_1^2}{(2-\alpha)^2} - (\mu_1-\mu_2)^2(1-\alpha) + (\sigma_1^2-\sigma_2^2)\alpha\right),$$

which is negative when $|\mu_1-\mu_2| \geq \frac{\sigma_1}{2}$, then positive when α increases. Therefore, the second term in Eq. (48) is first decreasing in α then increasing in α. Thus, there are two cases for the minimization in Eq. (48), either the maximum of the MSE's occurs when the two terms in Eq. (48) are equal, or the MSE occurs at the minimum of the second term.

Case 1: Switch and no-switch error is equal. We first consider the case that we need to balance $\left(\frac{\alpha}{2-\alpha}\right)\sigma_1^2$ and $\left(\frac{\alpha}{2-\alpha}\right)\sigma_1^2+(1-\alpha)^2(\mu_1-\mu_2)^2+\alpha^2(\sigma_2^2-\sigma_1^2)$. Solving the equation

$$\left(\frac{\alpha}{2-\alpha}\right)\sigma_1^2 = \left(\frac{\alpha}{2-\alpha}\right)\sigma_1^2 + (1-\alpha)^2(\mu_1-\mu_2)^2 + \alpha^2(\sigma_2^2-\sigma_1^2),$$

we obtain that the unique optimal α is

$$\alpha = \frac{|\mu_1-\mu_2|}{|\mu_1-\mu_2| + \sqrt{\sigma_1^2-\sigma_2^2}}.$$

Case 2: Minimum of switch error Next, we consider the case where switch error is minimized. Note that for $\alpha \in [0,1]$, the term $\frac{\sigma_1^2}{(2-\alpha)^2}$ is increasing and convex in α, and the term $-(\mu_1-\mu_2)^2(1-\alpha)+(\sigma_2^2-\sigma_1^2)\alpha$ is increasing and linear in α. When $\alpha = 0$, we have

$$2\left(\frac{\sigma_1^2}{(2-\alpha)^2} - (\mu_1-\mu_2)^2(1-\alpha) + (\sigma_2^2-\sigma_1^2)\alpha\right) = 2\left(\frac{\sigma_1^2}{4} - (\mu_1-\mu_2)^2\right) \leq 0$$

since $|\mu_1-\mu_2| \geq \frac{\sigma_1}{2}$ by assumption. Further, when $\alpha = 1$,

$$2\left(\frac{\sigma_1^2}{(2-\alpha)^2} - (\mu_1-\mu_2)^2(1-\alpha) + (\sigma_2^2-\sigma_1^2)\alpha\right) = \sigma_2^2 \geq 0.$$

Therefore, there is a unique solution for

$$2\left(\frac{\sigma_1^2}{(2-\alpha)^2} - (\mu_1-\mu_2)^2(1-\alpha) + (\sigma_2^2-\sigma_1^2)\alpha\right) = 0,$$

when $\alpha \in [0,1]$.

Finally, we need to identify the condition for when each these two cases is where the maximum MSE occurs. Let $\alpha^* = \frac{|\mu_1-\mu_2|}{|\mu_1-\mu_2|+\sqrt{\sigma_1^2-\sigma_2^2}}$. Since the derivative of the MSE after the service quality switch is $2\left(\frac{\sigma_1^2}{(2-\alpha)^2} - (\mu_1-\mu_2)^2(1-\alpha) + (\sigma_1^2-\sigma_2^2)\alpha\right)$, this error is first decreasing in α then increasing in α for $\alpha \in [0,1]$ and it is straight-forward to show it has a unique crossing point with the no-switch error, $\frac{\alpha\sigma_1^2}{(2-\alpha)}$, at α^*. Thus, when $2(\sigma_1^2/(2-\alpha^*)^2 - (\mu_1-\mu_2)^2(1-\alpha^*) + (\sigma_2^2-\sigma_1^2)\alpha^*) \leq 0$, the optimal α will be at the crossing point

$$\alpha = \alpha^*,$$

since the no-switch error is increasing and dominates the switch error for $\alpha \in (\alpha^*, 1]$. Otherwise, the optimal α will be the solution that minimizes the switching error,

$$2\left(\frac{\sigma_1^2}{(2-\alpha)^2} - (\mu_1-\mu_2)^2(1-\alpha) + (\sigma_2^2-\sigma_1^2)\alpha\right) = 0$$

for $\alpha \in [0,1]$. □

Proof of Lemma 8. Note that for an α-moving average system, the mean squared error at steady state is

$$\begin{aligned}
\mathbb{E}[(R_t-\mu)^2] &= \mathbb{E}\left[\left(\sum_{i=0}^{\infty}\alpha(1-\alpha)^i s_{t-i} - \mu\right)^2\right] \\
&= \mathbb{E}\left[\left(\sum_{i=0}^{\infty}\alpha(1-\alpha)^i (s_{t-i} - \mu)\right)^2\right] \\
&= \mathbb{E}\left[\sum_{i=1}^{\infty}\sum_{j=1}^{\infty}\left(\alpha(1-\alpha)^i(s_{t-i}-\mu)\right)\left(\alpha(1-\alpha)^j(s_{t-j}-\mu)\right)\right] \\
&= \sum_{i=1}^{\infty}\alpha^2(1-\alpha)^{2i}\mathbb{E}\left[(s_{t-i}-\mu)^2\right] \\
&= \sum_{i=1}^{\infty}\alpha^2(1-\alpha)^{2i}\sigma^2 \\
&= \left(\frac{\alpha}{2-\alpha}\right)\sigma^2
\end{aligned} \tag{49}$$

where the first equality follows from the definition of α-moving average system, and the second equality follows from the fact that $\sum_{i=0}^{\infty}\alpha(1-\alpha)^i = 1$. Further, the fourth equality follows from the fact that $s_{t-i}-\mu$ and $s_{t-j}-\mu$ are independent when $i \neq j$, i.e., $\mathbb{E}[(s_{t-i}-\mu)(s_{t-j}-\mu)] = \mathbb{E}[(s_{t-i}-\mu)]\mathbb{E}[(s_{t-j}-\mu)] = 0$. The final equality follows from the fact that the sum of the geometric sequence $\alpha^2(1-\alpha)^{2i}$ is $\frac{\alpha}{2-\alpha}$. Therefore, for an α-moving average system is be c-consistent, we need

$$\left(\frac{\alpha}{2-\alpha}\right)\sigma^2 \leq c\sigma^2.$$

Solving the inequality above, we obtain that α is at most $\frac{2c}{1+c}$.

Similarly, the mean squared error of an L-sliding window rating system is

$$\begin{aligned}
\mathbb{E}[(R_t-\mu)^2] &= \mathbb{E}\left[\left(\sum_{i=0}^{L-1}\frac{s_{t-i}}{L}-\mu\right)^2\right] \\
&= \mathbb{E}\left[\left(\sum_{i=0}^{L-1}\frac{s_{t-i}-\mu}{L}\right)^2\right] \\
&= \sum_{i=0}^{L-1}\mathbb{E}\left[\left(\frac{s_{t-i}-\mu}{L}\right)^2\right] \\
&= \frac{\sigma^2}{L}.
\end{aligned}$$

As before, the third equality follows from $s_{t^*+1-i}-\mu$ is independent of $s_{t-j}-\mu$ for all $i \neq j$, all the other equalities follow from simplification. For a L-sliding window system is to be c-consistent, we need

$$\frac{\sigma^2}{L} \leq c\sigma^2.$$

Since L needs to be an integer, the smallest L is $L = \lceil \frac{1}{c} \rceil$. □

Proof of Theorem 11. First, since both the moving average and sliding window systems are exactly c-consistent, by Lemma 8 we can write their parameters as a function of c i.e. $\alpha = \frac{2c}{1+c}$and $L = \frac{1}{c}$. Now for the first two parts of this proof we will write down the MSE for two rating systems, then in part (c) we will compare them.

(a) First we consider the L-sliding window rating system immediately after the service quality switch occurs, i.e., when $s_i \sim F_1$ for $i \le t^*$ and $s_i \sim F_2$ for $i > t^*$. By definition, the MSE of the L-sliding window rating system at $t^* + 1$ is,

$$
\begin{aligned}
\mathbb{E}[(R_{t^*+1} - \mu_2)^2 | s_{t^*+1} \sim F_2] &= \mathbb{E}\left[\left(\frac{\sum_{i=0}^{L-1} s_{t^*+1-i}}{L} - \mu_2\right)^2\right] \\
&= \mathbb{E}\left[\left(\frac{\sum_{i=1}^{L-1}(s_{t^*+1-i} - \mu_1)}{L} + \frac{\sum_{i=1}^{L-1}(\mu_1 - \mu_2)}{L} + \frac{(s_{t^*+1} - \mu_2)}{L}\right)^2\right] \\
&= \frac{\sum_{i=1}^{L-1} \mathbb{E}(s_{t^*+1-i} - \mu_1)^2}{L^2} + \left(\frac{\sum_{i=1}^{L-1}(\mu_1 - \mu_2)}{L}\right)^2 + \frac{\mathbb{E}(s_{t^*+1} - \mu_2)^2}{L^2} \\
&= \frac{\sum_{i=1}^{L-1} \sigma_1^2}{L^2} + \left(1 - \frac{1}{L}\right)^2 (\mu_1 - \mu_2)^2 + \frac{\sigma_2^2}{L^2} \\
&= \frac{\sigma_1^2}{L} + \left(1 - \frac{1}{L}\right)^2 (\mu_1 - \mu_2)^2 + \frac{\sigma_2^2 - \sigma_1^2}{L^2} \\
&= c\sigma_1^2 + (1-c)^2(\mu_1 - \mu_2)^2 + c^2(\sigma_2^2 - \sigma_1^2),
\end{aligned}
$$

where the first equality follows from the definition of the L-sliding window rating system. The second equality follows from decomposing $\mu_2 = \sum_{i=0}^{L-1} \frac{\mu_2}{L}$ in a way similar to the proof of Theorem 10. The third equality follows from the fact that $s_{t^*+1-i} - \mu_1$ for $i \le t^*$ and $s_{t^*+1} - \mu_2$ are independent with each other for $i \ge 1$ and in expectation each term is 0 i.e., $E[(s_{t^*+1-i} - mu_1)] = 0$ and $E[(s_{t^*+1} - \mu_2)] = 0$. Finally, the remaining equations follow from simplification and reorganization, as well as plugging in that $L = \frac{1}{c}$ which gives the desired expression.

(b) Next, we consider the α-moving average rating system immediately after the service quality switch occurs. In this case, the MSE is,

$$
\begin{aligned}
&E[(R_{t^*+1}-\mu_2)^2|s_{t^*+1}\sim F_2] \\
&=\mathbb{E}\left[\left(\sum_{i=0}^{\infty}\alpha(1-\alpha)^i s_{t-i}-\mu_2\right)^2\right] \\
&=\mathbb{E}\left[\left(\sum_{i=1}^{\infty}\alpha(1-\alpha)^i(s_{t^*+1-i}-\mu_1)+\sum_{i=1}^{\infty}\alpha(1-\alpha)^i(\mu_1-\mu_2)+\alpha(s_{t^*+1}-\mu_2)\right)^2\right] \\
&=\sum_{i=1}^{\infty}\alpha^2(1-\alpha)^{2i}\mathbb{E}[(s_{t^*+1-i}-\mu_1)^2]+\left(\sum_{i=1}^{\infty}\alpha(1-\alpha)^i(\mu_1-\mu_2)\right)^2+\alpha^2\mathbb{E}[(s_{t^*+1}-\mu_2)^2] \\
&=\sum_{i=1}^{\infty}\alpha^2(1-\alpha)^{2i}\sigma_1^2+((1-\alpha)(\mu_1-\mu_2))^2+\alpha^2\sigma_2^2 \\
&=\sum_{i=0}^{\infty}\alpha^2(1-\alpha)^{2i}\sigma_1^2+(1-\alpha)^2(\mu_1-\mu_2)^2+\alpha^2\left(\sigma_2^2-\sigma_1^2\right) \\
&=c\sigma_1^2+\frac{(1-c)^2}{(1+c)^2}(\mu_1-\mu_2)^2+\frac{c^2}{\left(\frac{1}{2}+\frac{c}{2}\right)^2}(\sigma_2^2-\sigma_1^2),
\end{aligned}
$$

where the first equality follows from the definition of the α-moving average rating system that is operating in steady-state. The second equality follows from the decomposition of μ_2 as above. The third equality follows from the independence of $s_{t^*+1-i}-\mu_1$ for $i\geq 1$ and $s_{t^*+1}-\mu_2$ and that each terms expected value is 0. Finally, the remaining equalities follow from simplification and plugging in $\alpha=\frac{2c}{1+c}$ and $1-\alpha=\frac{1-c}{1+c}$ to give the desired expression.

(c) Now armed with parts (a) and (b) we are ready to make a comparison between the two systems. Specifically, we will compare their maximum steady-state error. Recall from Theorem 10, for rating systems in steady-state the maximum expected error occurs either before the switch, where the error is due to variance in the system, or immediately after the switch. Since both systems are exactly c-consistent, their expected errors before the switch are identical. Thus to compare their maximum error, we need only look at which incurs more error after the switch which we can do by comparing the expressions in (a) and (b).

Now note that when $\sigma_1^2\geq\sigma_2^2$, we always have $\sigma_2^2-\sigma_1^2\leq 0$. Further note $\frac{(1-c)^2}{(1+c)^2}\leq(1-c)^2$ and $\frac{c^2}{\left(\frac{1}{2}+\frac{c}{2}\right)^2}\geq c^2$ for any $c\in[0,1]$ (recall c is always less than or equal to 1 by the definition of c-consistency). By applying these two inequalities to the MSE after the switch for the L-sliding window system we obtain,

$$c\sigma_1^2 + \frac{(1-c)^2}{(1+c)^2}\left(\mu_1-\mu_2\right)^2 + \frac{c^2}{\left(\frac{1}{2}+\frac{c}{2}\right)^2}(\sigma_2^2-\sigma_1^2) \leq c\sigma_1^2 + (1-c)^2(\mu_1-\mu_2)^2 + c^2(\sigma_2^2-\sigma_1^2),$$

which is the MSE after the service quality switch for the α-moving average system. Thus, the maximum MSE of the α-moving average rating system is always less than the maximum MSE of the L-sliding window rating system when $\sigma_1^2 \geq \sigma_2^2$. □

C.2.2 Omitted Proofs from Section 4.3

Proof of Proposition 1. In this proposition, we will solve the Stackelberg game described in Section 4.3. To do so, fix penalty term β and fee f, and consider some service provider with service distribution F_q such that $\mathbb{E}[F_q] = q$. Suppose the service provider chooses to transact on the platform with probability p. First, we will solve for the asymptotic expected rating:

$$\begin{aligned}
\lim_{t\to\infty} \mathbb{E}[R_t] = \mathbb{E}[\lim_{t\to\infty} R_t] &:= \mathbb{E}[R_\infty] \\
&= (1-\alpha)\mathbb{E}[R_\infty] + \alpha\left((1-ph\left(\mathbb{E}[R_\infty]\right))\beta + ph\left(\mathbb{E}[R_\infty]\right)q\right) \\
&= (1-\alpha)\mathbb{E}[R_\infty] + \alpha\left((1-p\left(a+(b-a)\mathbb{E}[R_\infty]\right))\beta + p\left(a+(b-a)\mathbb{E}[R_\infty]\right)q\right) \\
&= \frac{\beta + p(q-\beta)a}{1-p(q-\beta)(b-a)},
\end{aligned}$$

where the interchange in the first equality follows by dominated convergence where R_∞ is the limiting distribution of ratings for the service provider, in the second line equality we plugged in Eq. (17), in the third line we used the assumption that h is linear, and in the fourth line we simplify.

Thus in limit the platform and provider earn average profit per time period of:

$$\begin{aligned}
\text{Platform Profit} &:= rfp\mathbb{E}[h(R_\infty)] = frp\left(a + (b-a)\left(\frac{\beta+p(q-\beta)a}{1-p(q-\beta)(b-a)}\right)\right) \\
&= rfp\left(\frac{a+\beta(b-a)}{1-p(q-\beta)(b-a)}\right), \\
\text{Provider Profit} &:= r(1-f)p\mathbb{E}[h(R_\infty)] + r(1-p)\mathbb{E}[h(R_\infty)] \\
&= r(1-fp)\mathbb{E}[h(R_\infty)] \\
&= r(1-fp)\left(\frac{a+\beta(b-a)}{1-p(q-\beta)(b-a)}\right), \qquad (50)
\end{aligned}$$

where recall that p $\mathbb{E}[h(R_\infty)]$ is the probability the service provider is solicited ($\mathbb{E}[h(R_\infty)]$) and transacts on the platform (p) in steady state. Given these expressions, we can now solve the Stackelberg game where first the platform (the leader) sets f and β, then the provider (the follower) sets p. The derivative in p of the provider profit is,

$$\frac{d}{dp} r(1 - fp) \left(\frac{a + \beta(b-a)}{1 - p(q-\beta)(b-a)} \right) = r \frac{(a(1-\beta) + b\beta) \left((q-\beta)(b-a) - f \right)}{\left(1 - p(q-\beta)(b-a)\right)^2},$$

which is positive as long as $q \geq \beta + f/(b-a)$. If the derivative is positive, then note the provider profit is increasing in p, and thus it is optimal for the provider to set p to its maximum value of 1 and always faithfully transact on the platform. Otherwise, p should be pushed to 0, so the provider eventually will always disintermediate, and the profit for the platform from this provider goes to 0. This implies that for the platform the optimal fee, f^* is such that $f^* = (q - \beta)(b - a)$. □

C.2.3 Omitted Proofs from Section 4.4

Proof of Theorem 12. Fix some distribution of quality $q \sim G$ supported on $[\tau, 1]$, and some linear link function $h(x) = a + (b - a)x$ where $0 \leq a < b \leq 1$. Suppose (f^*, β^*) are the distribution dependent profit-maximizing choice of parameters. By Proposition 1, for these choices of fee and penalty to be profit-maximizing they must satisfy the constraint that $\beta^* + f^*/(b - a) = \tau$, or equivalently, $f^* + h(\beta^*) = h(\tau)$ by definition of $h(\cdot)$[1]. We will refer to choices of f and β such that $f + h(\beta) = h(\tau)$ as *feasible*, and in this proof derive feasible choices of f and β which always approximate the optimal profit.

Reducing the formulation. Let $\gamma(f, \beta, q) := fr\left(\frac{a+\beta(b-a)}{1-(q-\beta)(b-a)}\right)$ which, for feasible f and β by Proposition 1, is the platform's equilibrium profit from a provider as a function of the fee (f), rating parameter (β), and providers expected quality (q). By definition of $h(\cdot)$ we can rewrite γ as,

$$\gamma(f, \beta, q) = fr\left(\frac{h(\beta)}{1 - (h(q) - h(\beta))}\right).$$

[1]Note that Proposition 1 states this constraint as an inequality. However, if some choice of (f, β) satisfies this constraint with strict inequality then the platform can strictly increase profit by increasing f to make the constraint hold with equality.

Note that for any fixed, feasible (β, f), and for any $q \in [\tau, 1]$, $h(q) \geq h(\tau)$ and we have $\partial_q \gamma(\beta, f, q) \geq 0$ and $\partial_q^2 \gamma(\beta, f, q) \geq 0$. Thus for feasible choices of f and β, $\gamma(f, \beta, q)$ is an increasing and convex function in q. Now we consider lower bounding the following ratio:

$$\min_{G \in \mathcal{D}} \frac{\mathbb{E}_{q \sim G}[\gamma(f, \beta, q)]}{\mathbb{E}_{q \sim G}[\gamma(f^*, \beta^*, q)]}, \tag{51}$$

where $\mathcal{D}$ is the class of distributions of mean service quality supported on $[\tau, 1]$, and where $(f^*, \beta^*) = \arg\max_{f,\beta} \mathbb{E}_{q \sim G}[\gamma(f, \beta, q)]$ such that $f^* + h(\beta^*) = h(\tau)$. We will show there exists specifications of f and β, and a function dependent only on τ denoted $\eta(\tau)$, such that for any optimal f^*, β^* and any q we have $\gamma(f, \beta, q)/\gamma(f^*, \beta^*, q) \geq \eta(\tau)$. This in turn implies that,

$$\min_{G \in \mathcal{D}} \frac{\mathbb{E}_{q \sim G}[\gamma(f, \beta, q)]}{\mathbb{E}_{q \sim G}[\gamma(f^*, \beta^*, q)]} \geq \eta(\tau).$$

To that end, note that the ratio equals:

$$\begin{aligned}
\frac{\gamma(f, \beta, q)}{\gamma(f^*, \beta^*, q)} &= \frac{fr\left(\frac{h(\beta)}{1-(h(q)-h(\beta))}\right)}{f^* r\left(\frac{h(\beta^*)}{1-(h(q)-h(\beta^*))}\right)} \\
&= \frac{f}{f^*}\frac{h(\beta)}{h(\beta^*)}\frac{1-h(q)+h(\beta^*)}{1-h(q)+h(\beta)} \\
&= \frac{f}{f^*}\frac{h(\beta)}{h(\beta^*)}\left(1 + \frac{h(\beta^*)-h(\beta)}{1+h(\beta)-h(q)}\right),
\end{aligned}$$

which is always a monotone function in q that is minimized at either $q = \tau$ or $q = 1$, depending on whether $\beta - \beta^* > 0$ or not. Thus, to lower bound the ratio in Eq. (51), we only need to bound the ratio $\gamma(f, \beta, q)/\gamma(f^*, \beta^*, q)$ at $q = \tau$ and $q = 1$.

Bounding the ratio. To bound the ratio we consider two cases, when $q = \tau$ and when $q = 1$. When $q = \tau$, the ratio can be simplified as

$$\frac{\gamma(f, \beta, \tau)}{\gamma(f^*, \beta^*, \tau)} = \frac{fr\left(\frac{h(\beta)}{1-(h(\tau)-h(\beta))}\right)}{f^* r\left(\frac{h(\beta^*)}{1-(h(\tau)-h(\beta^*))}\right)} = \frac{f\left(\frac{h(\tau)-f}{1-f}\right)}{f^*\left(\frac{h(\tau)-f^*}{1-f^*}\right)} \tag{52}$$

where the second equality follows from the constraint that $h(\beta) = h(\tau) - f$. Similarly, when $q = 1$, the ratio can be simplified as

$$\frac{\gamma(f,\beta,1)}{\gamma(f^*,\beta^*,1)} = \frac{fr\left(\frac{h(\beta)}{1-(h(1)-h(\beta))}\right)}{f^*r\left(\frac{h(\beta^*)}{1-(h(1)-h(\beta^*))}\right)} = \frac{f\left(\frac{h(\tau)-f}{1-h(1)+h(\tau)-f)}\right)}{f^*\left(\frac{h(\tau)-f^*}{1-h(1)+h(\tau)-f^*}\right)} \tag{53}$$

where again the second equality follows since $h(\beta) = h(\tau) - f$.

Now we aim to choose feasible f and $h(\beta)$[2] to lower bound the ratio of $\frac{\gamma(f,\beta,q)}{\gamma(f^*,\beta^*,q)}$ at $q = \tau$ and $q = 1$. Formally, we wish to solve:

$$\begin{aligned} \max_{f,h(\beta)} \min_{\tilde{f},h(\tilde{\beta})} & \left\{\frac{\gamma(f,\beta,\tau)}{\gamma(\tilde{f},\tilde{\beta},\tau)}, \frac{\gamma(f,\beta,1)}{\gamma(\tilde{f},\tilde{\beta},1)}\right\}, \\ \text{s.t. } & f + h(\beta) = \tilde{f} + h(\tilde{\beta}) = h(\tau) \end{aligned} \tag{54}$$

where note the unknown optimal parameters f^*, β^* are replaced by adversarially chosen $\tilde{f}$, $\tilde{\beta}$. Thus the solution to Eq. (54) lower bounds Eq. (51).

Now, for any feasible f, β the first term in the inner min of Eq. (54) is:

$$\begin{aligned} \min_{\tilde{f}_\tau,h(\tilde{\beta}_\tau)} \frac{\gamma(f,\beta,\tau)}{\gamma(\tilde{f}_\tau,\tilde{\beta}_\tau,\tau)} = \min_{\tilde{f}_\tau,h(\tilde{\beta}_\tau)} \frac{f\left(\frac{h(\tau)-f}{1-f}\right)}{\tilde{f}_\tau\left(\frac{h(\tau)-\tilde{f}_\tau}{1-\tilde{f}_\tau}\right)} &= \frac{f\left(\frac{h(\tau)-f}{1-f}\right)}{(1-\sqrt{1-h(\tau)})\left(\frac{h(\tau)-(1-\sqrt{1-h(\tau)})}{1-(1-\sqrt{1-h(\tau)})}\right)} \\ &= \frac{f\left(\frac{h(\tau)-f}{1-f}\right)}{h^2(\tau)/\left(1+\sqrt{1-h(\tau)}\right)^2} \end{aligned}$$

where the first equality follows by Eq. (52) and the second by noting the expression is minimized by $\tilde{f}_\tau = 1 - \sqrt{1-h(\tau)}$, $h(\tilde{\beta}_\tau) = h(\tau) - 1 + \sqrt{1-h(\tau)}$. Similarly, for any feasible f, β, the second term in Eq. (54) is:

$$\begin{aligned} &\min_{\tilde{f}_1,h(\tilde{\beta}_1)} \frac{\gamma(f,\beta,1)}{\gamma(\tilde{f},\tilde{\beta},1)} = \min_{\tilde{f}_1,h(\tilde{\beta}_1)} \frac{f\left(\frac{h(\tau)-f}{1-h(1)+h(\tau)-f)}\right)}{\tilde{f}_1\left(\frac{h(\tau)-\tilde{f}_1}{1-h(1)+h(\tau)-\tilde{f}_1}\right)} \\ &= \frac{f\left(\frac{h(\tau)-f}{1-h(1)+h(\tau)-f)}\right)}{\left((1-h(1))+h(\tau)-\sqrt{(1-h(1))(1-h(1)+h(\tau))}\right)\left(\frac{h(\tau)-((1-h(1))+h(\tau)-\sqrt{(1-h(1))(1-h(1)+h(\tau))})}{1-h(1)+h(\tau)-((1-h(1))+h(\tau)-\sqrt{(1-h(1))(1-h(1)+h(\tau))})}\right)} \\ &= \frac{f\left(\frac{h(\tau)-f}{1-h(1)+h(\tau)-f)}\right)}{h(\tau)-2\left(\sqrt{(1-h(1))(1-h(1)+h(\tau))}-(1-h(1))\right)} \end{aligned}$$

[2]As h is a linear and thus invertible function, the optimal choice of $h(\beta)$ implies the optimal choice of β. It will be more convenient to work $h(\beta)$ for the remainder of the proof.

where the first equality follows by Eq. (53) and the second by noting the expression is minimized by $\tilde{f}_1 = (1-h(1))+h(\tau)-\sqrt{(1-h(1))(1-h(1)+h(\tau))}$ and $h(\tilde{\beta}_1) = \sqrt{(1-h(1))(1-h(1)+h(\tau))}-(1-h(1))$.

The inner minimization of Eq. (54) is achieved at one of these pairs, either $(\tilde{f}_\tau, \tilde{\beta}_\tau)$ or $(\tilde{f}_1, \tilde{\beta}_1)$. Now we wish to calculate feasible choices of $(f, h(\beta))$ which maximize the minimum of these two options. However, directly solving this maximization problem is difficult. Instead, consider the following heuristic choices of f and $h(\beta)$,

$$\overline{f} = \frac{\left(2+\sqrt{1-h(\tau)}\right)h(\tau)}{2\left(1+\sqrt{1-h(\tau)}\right)}, \quad h(\overline{\beta}) = \frac{\left(\sqrt{1-h(\tau)}\right)h(\tau)}{2(1+\sqrt{1-h(\tau)})}$$

Note this choice of $\overline{f}$ and $h(\overline{\beta})$ are both positive and sum to $h(\tau)$, thus they are a feasible choice of parameters. With this heuristic choice of the fee, the first term of Eq. (54) becomes

$$\begin{aligned}\frac{\gamma(\overline{f}, \overline{\beta}, \tau)}{\gamma(\tilde{f}_\tau, \tilde{\beta}_\tau, \tau)} &= \frac{\overline{f}(h(\tau)-\overline{f})/(1-\overline{f})}{h^2(\tau)/\left(1+\sqrt{1-h(\tau)}\right)^2} \\ &= \frac{2+\sqrt{1-h(\tau)}}{2\left(1+\sqrt{1-h(\tau)}\right)} \\ &\geq \frac{3}{4}\end{aligned}$$

where the first equality follows from plugging in the values of $(\tilde{f}_\tau, \tilde{\beta}_\tau)$, the second equality follows plugging the heuristic choice of $\overline{f}$, and the inequality follows from minimizing over possible values of $\sqrt{1-h(\tau)}$.

We also notice that with the heuristic choice of f, the first ratio in Eq. (54) satisfies

$$\frac{\gamma(\overline{f}, \overline{\beta}, \tau)}{\gamma(\tilde{f}_\tau, \tilde{\beta}_\tau, \tau)} = \frac{2+\sqrt{1-h(\tau)}}{2\left(1+\sqrt{1-h(\tau)}\right)} = \frac{\overline{f}}{h(\tau)}.$$

Now, we aim to show that $\frac{\overline{f}}{h(\tau)}$ is a lower bound of the second ratio in Eq. (54), i.e.,

$$\frac{\gamma(\overline{f}, \overline{\beta}, 1)}{\gamma(\tilde{f}_1, \tilde{\beta}_1, 1)} \geq \frac{f}{h(\tau)}.$$

Specifically, we consider the following equation,

$$\frac{\gamma(\overline{f},\overline{\beta},1)/\overline{f}}{\gamma(\tilde{f}_1,\tilde{\beta}_1,1)/h(\tau)} = \frac{h(\tau)\left(h(\tau)-1+\sqrt{1-h(\tau)}\right)}{\left(1+\sqrt{1-h(\tau)}+h(\tau)-2h(1)\right)\left(h(\tau)+2(1-h(1)-\sqrt{(1-h(1))(1-h(1)+h(\tau))})\right)},$$

where the equality follows from plugging in $\tilde{f}_1, \tilde{\beta}_1$ and the heuristic choice $\overline{f}$. Finally, by a lengthy purely algebraic calculation we can verify that for any $h(\tau) \le h(1) \le 1$, we always have

$$\frac{h(\tau)\left(h(\tau)-1+\sqrt{1-h(\tau)}\right)}{\left(1+\sqrt{1-h(\tau)}+h(\tau)-2h(1)\right)\left(h(\tau)+2(1-h(1)-\sqrt{(1-h(1))(1-h(1)+h(\tau))})\right)} \ge 1.$$

In the interest of brevity we refer readers to `https://github.com/tcui-pitt/Ratings/Thm3.nb` for detailed computations using a computer algebra system.

Finally, combining the results above, we conclude that

$$\text{Eq. (51)} \ge \text{Eq. (54)} \ge \frac{2+\sqrt{1-h(\tau)}}{2\left(1+\sqrt{1-h(\tau)}\right)} \ge \frac{3}{4}.$$

Therefore, the choice of $f = \frac{\left(2+\sqrt{1-h(\tau)}\right)h(\tau)}{2\left(1+\sqrt{1-h(\tau)}\right)}$ always earns at least $\frac{3}{4}^{th}$ of the optimal profit. Further, this ratio is also tight for a corresponding class of point mass distributions. For a complete construction of the tight example see Example 9. □

C.3 Omitted Figures

In this section we include additional motivating figures that were omitted from main body.

C.3.1 Stale Rating Systems in Practice

In this subsection we give some examples of ratings on the platforms Yelp and Google Reviews. We include examples both sequentially and temporally stale ratings on both platforms.

Figure 29: Examples of Popular Rating Systems.

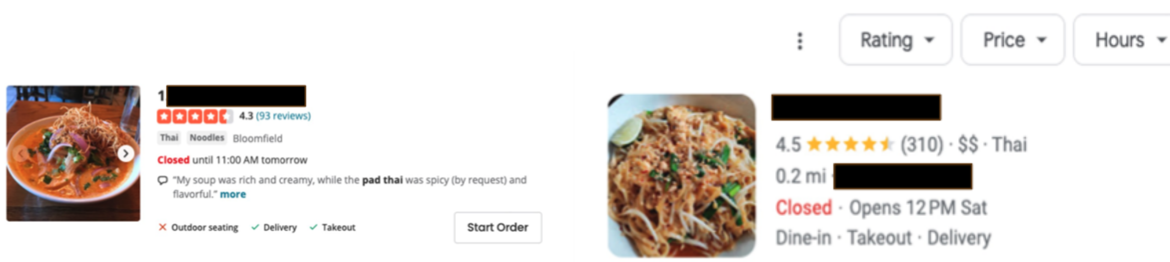

(a) Yelp.com Rating (b) Google Rating

Note. Depicted are ratings for a local Thai restaurant on `yelp.com` (left) and `google.com` (right). Note in each, both a numeric rating and the number of reviews that inform the rating are clearly displayed. The individual reviews that form each rating are also available.

Figure 30: Sequentially Stale Reviews in Practice.

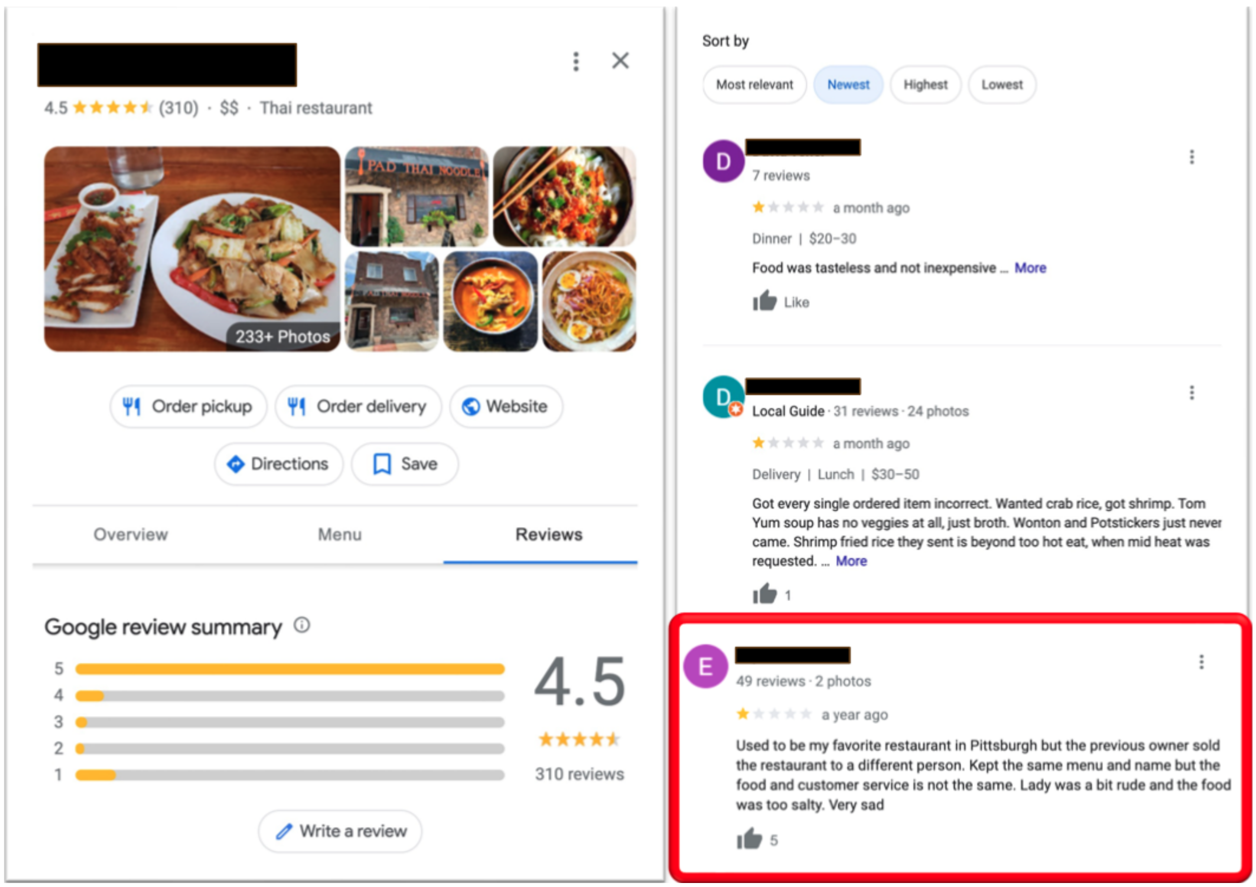

Note. Depicted is the rating and a review summary for a local restaurant (left), juxtaposed with some recent low quality reviews (right). Although the rating is currently high (4.5 out of 5), it is sequentially stale, that is substantially informed by older reviews which evidently may not represent the current service quality. The highlighted review on the bottom left gives one possible explanation for the abrupt change in underlying quality: there was a recent change in management.

Figure 31: Temporally Stale Reviews in Practice.

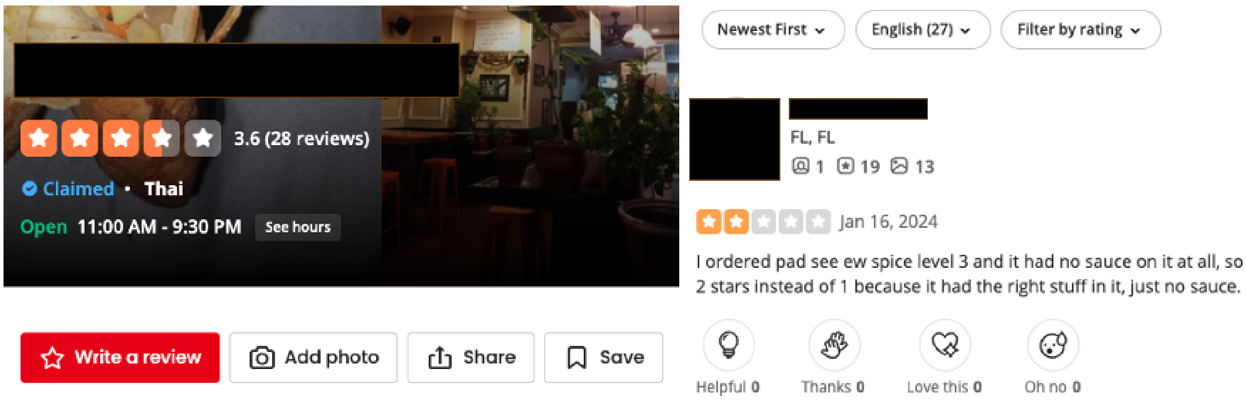

Note. Depicted is the rating and review summary for a local Thai restaurant, captured on May 19th, 2024. The rating is of moderate quality (enough to show up in the first page of yelp recommendations for Thai food) however the most recent review is more than five months old having been written in January, 2024. Indeed, this temporally stale rating is quite inaccurate since this restaurant closed on March of 2024.

Bibliography

[1] Takanori Adachi and Mark J Tremblay. Bargaining in two-sided markets. *Available at SSRN*, 3388383, 2019.

[2] Afp. Singapore woos singles in latest dating campaign. August 2010.

[3] Vipul Agrawal and Sridhar Seshadri. Risk intermediation in supply chains. *Iie Transactions*, 32(9):819–831, 2000.

[4] Vishal V. Agrawal, Mark Ferguson, L. Beril Toktay, and Valerie M. Thomas. Is leasing greener than selling? *Management Science*, 58(3):523–533, 2012.

[5] AKITA. Japan's state-owned version of tinder. October 2019.

[6] Gad Allon, Achal Bassamboo, and Eren B Cil. Large-scale service marketplaces: The role of the moderating firm. *Management Science*, 58(10):1854–1872, 2012.

[7] Samaneh Aminikhanghahi and Diane J Cook. A survey of methods for time series change point detection. *Knowledge and information systems*, 51(2):339–367, 2017.

[8] Ali Aouad, Adam N Elmachtoub, Kris J Ferreira, and Ryan McNellis. Market segmentation trees. *arXiv preprint arXiv:1906.01174*, 2019.

[9] Christina Aperjis and Ramesh Johari. Designing aggregation mechanisms for reputation systems in online marketplaces. *Acm Sigecom Exchanges*, 9(1):1–4, 2010.

[10] Christina Aperjis and Ramesh Johari. Optimal windows for aggregating ratings in electronic marketplaces. *Management Science*, 56(5):864–880, 2010.

[11] Nick Arnosti, Ramesh Johari, and Yash Kanoria. Managing congestion in matching markets. *Manufacturing & Service Operations Management*, 23(3):620–636, 2021.

[12] Henry Assael and A Marvin Roscoe Jr. Approaches to market segmentation analysis. *Journal of Marketing*, 40(4):67–76, 1976.

[13] Susan Athey and Guido Imbens. Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27):7353–7360, 2016.

[14] Mark Bagnoli and Ted Bergstrom. Log-concave probability and its applications. *Economic theory*, 26(2):445–469, 2005.

[15] Sridhar Balasubramanian, Shantanu Bhattacharya, and Vish V. Krishnan. Pricing information goods: A strategic analysis of the selling and pay-per-use mechanisms. *Marketing Science*, 34(2):218–234, 2015.

[16] Santiago R Balseiro, Omar Besbes, and Gabriel Y Weintraub. Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Science*, 61(4):864–884, 2015.

[17] Amit Basu, Sreekumar Bhaskaran, and Rajiv Mukherjee. An analysis of search and authentication strategies for online matching platforms. *Management Science*, 65(5):2412–2431, 2019.

[18] Elena Belavina and Karan Girotra. The relational advantages of intermediation. *Management Science*, 58(9):1614–1631, 2012.

[19] Elena Belavina, Karan Girotra, Ken Moon, and Jiding Zhang. Matching in labor marketplaces: The role of experiential information. *Available at SSRN 3543906*, 2020.

[20] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The limits of price discrimination. *The American Economic Review*, 105(3):921–957, 2015.

[21] Dimitris Bertsimas, Jack Dunn, and Nishanth Mundru. Optimal prescriptive trees. *INFORMS Journal on Optimization*, 1(2):164–183, 2019.

[22] Omar Besbes and Ilan Lobel. Intertemporal price discrimination: Structure and computation of optimal policies. *Management Science*, 61(1):92–110, 2015.

[23] Purnima Bholowalia and Arvind Kumar. Ebk-means: A clustering technique based on elbow method and k-means in wsn. *International Journal of Computer Applications*, 105(9), 2014.

[24] Max Biggs, Wei Sun, and Markus Ettl. Model distillation for revenue optimization: Interpretable personalized pricing. In *International Conference on Machine Learning*, pages 946–956. PMLR, 2021.

[25] Kostas Bimpikis, Wedad J Elmaghraby, Ken Moon, and Wenchang Zhang. Managing market thickness in online business-to-business markets. *Management Science*, 66(12):5783–5822, 2020.

[26] Mike Blumenthal. Yelp's hidden headwinds: Declining user traffic and aging reviews. Technical report, 2022.

[27] Alex Bocknek. Swipe dating apps: Tinder, bumble, & more. March 2022.

[28] Luigi Buzzacchi and Tommaso M Valletti. Strategic price discrimination in compulsory insurance markets. *The Geneva Risk and Insurance Review*, 30(1):71–97, 2005.

[29] Gérard P Cachon and Pnina Feldman. Pricing services subject to congestion: Charge per-use fees or sell subscriptions? *Manufacturing & Service Operations Management*, 13(2):244–260, 2011.

[30] Trudy A Cameron and Michelle D James. Estimating willingness to pay from survey data: an alternative pre-test-market evaluation procedure. *Journal of Marketing Research*, 24(4):389–395, 1987.

[31] Marek Capinski and Tomasz Zastawniak. Mathematics for finance. *An Introduction*, pages 118–124, 2003.

[32] Christoph Carnehl, André Stenzel, and Peter Schmidt. Pricing for the stars: Dynamic pricing in the presence of rating systems. *Management Science*, 2023.

[33] Fangruo Chen. Market segmentation, advanced demand information, and supply chain performance. *Manufacturing & Service Operations Management*, 3(1):53–67, 2001.

[34] Mingliu Chen, Adam N Elmachtoub, and Xiao Lei. Matchmaking strategies for maximizing player engagement in video games. *Available at SSRN 3928966*, 2021.

[35] Xi Chen, Zachary Owen, Clark Pixton, and David Simchi-Levi. A statistical learning approach to personalization in revenue management. *Available at SSRN: https://ssrn.com/abstract=2579462*, 2015.

[36] Yuxin Chen, Sridhar Moorthy, and Z John Zhang. Research note-price discrimination after the purchase: Rebates as state-dependent discounts. *Management Science*, 51(7):1131–1140, 2005.

[37] Zhengxing Chen, Su Xue, John Kolen, Navid Aghdaie, Kazi A Zaman, Yizhou Sun, and Magy Seif El-Nasr. Eomm: An engagement optimized matchmaking framework. In *Proceedings of the 26th International Conference on World Wide Web*, pages 1143–1150, 2017.

[38] Henry J Claycamp and William F Massy. A theory of market segmentation. *Journal of Marketing Research*, 5(4):388–394, 1968.

[39] Maxime C Cohen, Ngai-Hang Zachary Leung, Kiran Panchamgam, Georgia Perakis, and Anthony Smith. The impact of linear optimization on promotion planning. *Operations Research*, 65(2):446–468, 2017.

[40] Maxime C. Cohen, Ilan Lobel, and Renato Paes Leme. Feature-based dynamic pricing. *Available at SSRN: https://ssrn.com/abstract=2737045*, 2016.

[41] Pascal Courty and Mario Pagliero. The impact of price discrimination on revenue: Evidence from the concert industry. *Review of Economics and Statistics*, 94(1):359–369, 2012.

[42] Simon Cowan. Welfare-increasing third-degree price discrimination. *The RAND Journal of Economics*, 47(2):326–340, 2016.

[43] Qilin Deng, Hao Li, Kai Wang, Zhipeng Hu, Runze Wu, Linxia Gong, Jianrong Tao, Changjie Fan, and Peng Cui. Globally optimized matchmaking in online games. In

Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, pages 2753–2763, 2021.

[44] Alexandre Dolgui and Jean-Marie Proth. Pricing strategies and models. *Annual Reviews in Control*, 34(1):101–110, 2010.

[45] Jean-Pierre Dubé, Zheng Fang, Nathan Fong, and Xueming Luo. Competitive price targeting with smartphone coupons. *Marketing Science*, 36(6):944–975, 2017.

[46] Anthony Dukes and Lin Liu. Online shopping intermediaries: The strategic design of search environments. *Management Science*, 62(4):1064–1077, 2016.

[47] Rosie Dunford, Quanrong Su, and Ekraj Tamang. The pareto principle. 2014.

[48] Glenn Ellison and Sara Fisher Ellison. Search, obfuscation, and price elasticities on the internet. *Econometrica*, 77(2):427–452, 2009.

[49] Glenn Ellison and Alexander Wolitzky. A search cost model of obfuscation. *The RAND Journal of Economics*, 43(3):417–441, 2012.

[50] Adam Elmachtoub, Jason Cheuk Nam Liang, and Ryan McNellis. Decision trees for decision-making under the predict-then-optimize framework. In *International Conference on Machine Learning*, pages 2858–2867. PMLR, 2020.

[51] Adam N Elmachtoub and Paul Grigas. Smart "predict, then optimize". *Management Science*, 68(1):9–26, 2022.

[52] Adam N Elmachtoub, Vishal Gupta, and Michael L Hamilton. The value of personalized pricing. *Management Science*, 2021.

[53] Adam N Elmachtoub and Michael L Hamilton. The power of opaque products in pricing. *Management Science*, 67(8):4686–4702, 2021.

[54] Ming Fan, Yong Tan, and Andrew B Whinston. Evaluation and design of online cooperative feedback mechanisms for reputation management. *IEEE Transactions on Knowledge and Data Engineering*, 17(2):244–254, 2005.

[55] Dave Fedewa, Chauncey Holder, Warren Teichner, and Bill Wiseman. Five-star growth: Using online ratings to design better products — mckinsey.com. `https://www.mckinsey.com/industries/consumer-packaged-goods/our-insights/five-star-growth-using-online-ratings-to-design-better-products`, 2021. [Accessed 12-04-2024].

[56] Kris Johnson Ferreira, Bin Hong Alex Lee, and David Simchi-Levi. Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & service operations management*, 18(1):69–88, 2016.

[57] Nikhil Garg and Ramesh Johari. Designing optimal binary rating systems. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1930–1939. PMLR, 2019.

[58] Evan Michael Gilbert. Antitrust and commitment issues: Monopolization of the dating app industry. *NYUL Rev.*, 94:862, 2019.

[59] Google. Reviews - google business profile help. Technical report, 2024.

[60] John C Gower. Some distance properties of latent root and vector methods used in multivariate analysis. *Biometrika*, 53(3-4):325–338, 1966.

[61] John C Gower. A comparison of some methods of cluster analysis. *Biometrics*, pages 623–637, 1967.

[62] Allan Gronlund, Kasper Larsen, Alexander Mathiasen, Jesper Nielsen, Stefan Schneider, and Mingzhou Song. Fast exact k-means, k-medians and bregman divergence clustering in 1d. *ArXiv e-prints: arXiv:1701.07204*, 2017.

[63] Andrei Hagiu and Julian Wright. Marketplace leakage. 2022.

[64] Hanna Halaburda, Mikołaj Jan Piskorski, and Pınar Yıldırım. Competing by restricting choice: The case of matching platforms. *Management Science*, 64(8):3574–3594, 2018.

[65] Sergiu Hart and Philip J Reny. Maximal revenue with multiple goods: Nonmonotonicity and other observations. *Theoretical Economics*, 10(3):893–922, 2015.

[66] Eryn Juan He, Sergei Savin, Joel Goh, and Chung-Piaw Teo. Off-platform threats in on-demand services. *Manufacturing & Service Operations Management*, 2023.

[67] Sherry He, Brett Hollenbeck, and Davide Proserpio. The market for fake reviews. *Marketing Science*, 41(5):896–921, 2022.

[68] Hugo Hopenhayn and Maryam Saeedi. Optimal ratings and market outcomes. Technical report, National Bureau of Economic Research, 2019.

[69] Johannes Hörner and Nicolas S Lambert. Motivational ratings. *The Review of Economic Studies*, 88(4):1892–1935, 2021.

[70] Jiali Huang, Ankur Mani, and Zizhuo Wang. The value of price discrimination in large random networks. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 243–244, 2019.

[71] Yan Huang, Stefanus Jasin, and Puneet Manchanda. "level up": Leveraging skill and engagement to maximize player game-play in online video games. *Information Systems Research*, 30(3):927–947, 2019.

[72] IBISWorld. Dating services in the us. June 2021.

[73] Krishnamurthy Iyer, Ramesh Johari, and Mukund Sundararajan. Mean field equilibria of dynamic auctions with learning. *Management Science*, 60(12):2949–2970, 2014.

[74] Adel Javanmard and Hamid Nazerzadeh. Dynamic pricing in high-dimensions. *Available at SSRN: https://ssrn.com/abstract=2855843*, 2016.

[75] Kinshuk Jerath, Serguei Netessine, and Senthil K Veeraraghavan. Revenue management with strategic customers: Last-minute selling and opaque selling. *Management Science*, 56(3):430–448, 2010.

[76] Ramesh Johari, Bar Light, and Gabriel Weintraub. Quality selection in two-sided markets: A constrained price discrimination approach. *arXiv preprint arXiv:1912.02251*, 2019.

[77] Nathan Kallus. Recursive partitioning for personalization using observational data. In *International conference on machine learning*, pages 1789–1798. PMLR, 2017.

[78] Yash Kanoria and Daniela Saban. Facilitating the search for partners on matching platforms. *Management Science*, 67(10):5990–6029, 2021.

[79] Nil Karacaoglu, Simin Li, and Ioannis Stamatopoulos. Disintermediation evidence from a cleaning platform. *Available at SSRN*, 2022.

[80] Michael L Katz. Price discrimination and monopolistic competition. *Econometrica: Journal of the Econometric Society*, pages 1453–1471, 1984.

[81] Jon Kleinberg, Christos Papadimitriou, and Prabhakar Raghavan. A microeconomic view of data mining. *Data mining and knowledge discovery*, 2(4):311–324, 1998.

[82] Jon Kleinberg, Christos Papadimitriou, and Prabhakar Raghavan. Segmentation problems. *Journal of the ACM (JACM)*, 51(2):263–280, 2004.

[83] Sergey Kovbasyuk and Giancarlo Spagnolo. Memory and Markets. *The Review of Economic Studies*, page rdad067, 06 2023.

[84] Andreas Krause and Daniel Golovin. Submodular function maximization. *Tractability*, 3:71–104, 2014.

[85] Konstantinos Ladas, Stylianos Kavadias, and Christoph Loch. Product selling vs. pay-per-use service: A strategic analysis of competing business models. *Management Science*, 0(0):null, 2021.

[86] Sonja Lehmann and Peter Buxmann. Pricing strategies of software vendors. *Business & Information Systems Engineering*, 1(6):452–462, 2009.

[87] Shujing Li and Jiaping Qiu. Financial product differentiation over the state space in the mutual fund industry. *Management Science*, 60(2):508–520, 2014.

[88] Jinan Lin, Tingting Nian, and Natasha Z Foutz. Disintermediation and its mitigation in online two-sided platforms: Evidence from airbnb. 2022.

[89] Ying Liu, Sudha Ram, Robert F Lusch, and Michael Brusco. Multicriterion market segmentation: a new model, implementation, and evaluation. *Marketing Science*, 29(5):880–894, 2010.

[90] Fei Long and Yunchuan Liu. Fake sales and ranking algorithms in online retail marketplace with sponsored advertising. *Available at SSRN 4009128*, 2022.

[91] Cynthia A Lowry and Douglas C Montgomery. A review of multivariate control charts. *IIE transactions*, 27(6):800–810, 1995.

[92] Cynthia A Lowry, William H Woodall, Charles W Champ, and Steven E Rigdon. A multivariate exponentially weighted moving average control chart. *Technometrics*, 34(1):46–53, 1992.

[93] Michael Luca. Designing online marketplaces: Trust and reputation mechanisms. *Innovation Policy and the Economy*, 17(1):77–93, 2017.

[94] Michael Luca and Georgios Zervas. Fake it till you make it: Reputation, competition, and yelp review fraud. *Management Science*, 62(12):3412–3427, 2016.

[95] Louise Matsakis. What Do Amazon's Star Ratings Really Mean? — wired.com. `https://www.wired.com/story/amazon-stars-ratings-calculated/`, May 2019.

[96] Mary Meisenzahl. These charts from match group show more people are turning to online dating during the pandemic. August 2020.

[97] Mohamed Mostagir and James Siderius. Strategic reviews. *Management Science*, 69(2):904–921, 2023.

[98] Susan M Mudambi and David Schuff. Research note: What makes a helpful online review? a study of customer reviews on amazon. com. *MIS quarterly*, pages 185–200, 2010.

[99] Chakravarthi Narasimhan. A price discrimination theory of coupons. *Marketing Science*, 3(2):128–147, 1984.

[100] George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. An analysis of approximations for maximizing submodular set functions—i. *Mathematical programming*, 14(1):265–294, 1978.

[101] Marius F. Niculescu and D. J. Wu. Economics of free under perpetual licensing: Implications for the software industry. *Information Systems Research*, 25(1):173–199, 2014.

[102] Chaoyue Niu, Zhenzhe Zheng, Fan Wu, Shaojie Tang, and Guihai Chen. Online pricing with reserve price constraint for personal data markets. *IEEE Transactions on Knowledge and Data Engineering*, 2020.

[103] Kari Paul. Uber to ban riders with low ratings: will you pass the test? Technical report, June 2019.

[104] Catherine Pearson. 'a decade of fruitless searching': The toll of dating app burnout. Aug 2022.

[105] Sheng Qiang and Mohsen Bayati. Dynamic pricing with demand covariates. *Available at SSRN: https://ssrn.com/abstract=2765257*, 2016.

[106] Alan P Reynolds, Graeme Richards, Beatriz de la Iglesia, and Victor J Rayward-Smith. Clustering rules: a comparison of partitioning and hierarchical clustering algorithms. *Journal of Mathematical Modelling and Algorithms*, 5(4):475–504, 2006.

[107] Ignacio Rios, Daniela Saban, and Fanyin Zheng. Improving match rates in dating markets through assortment optimization. *Manufacturing & Service Operations Management*, 25(4):1304–1323, 2023.

[108] Bob Ritchie and Clare Brindley. Disintermediation, disintegration and risk in the sme global supply chain. *Management Decision*, 38(8):575–583, 2000.

[109] PB Robinson and T Yr Ho. Average run lengths of geometric moving average charts by numerical methods. *Technometrics*, 20(1):85–93, 1978.

[110] Jean-Charles Rochet and Jean Tirole. Two-sided markets: a progress report. *The RAND journal of economics*, 37(3):645–667, 2006.

[111] Maryam Saeedi and Ali Shourideh. Optimal rating design under moral hazard. *arXiv preprint arXiv:2008.09529*, 2020.

[112] Adrien Saumard and Jon A Wellner. Log-concavity and strong log-concavity: a review. *Statistics surveys*, 8:45, 2014.

[113] Richard Schmalensee. Output and welfare implications of monopolistic third-degree price discrimination. *The American Economic Review*, 71(1):242–247, 1981.

[114] Erich Schubert and Peter J Rousseeuw. Faster k-medoids clustering: improving the pam, clara, and clarans algorithms. In *International conference on similarity search and applications*, pages 171–187. Springer, 2019.

[115] Erich Schubert and Peter J Rousseeuw. Fast and eager k-medoids clustering: O (k) runtime improvement of the pam, clara, and clarans algorithms. *Information Systems*, 101:101804, 2021.

[116] Shreyas Sekar and Auyon Siddiq. Platform disintermediation: Information effects and pricing remedies. *Available at SSRN 4378501*, 2023.

[117] Alex Shashkevich. Meeting online has become the most popular way u.s. couples connect, stanford sociologist finds. August 2019.

[118] Zijun Shi, Kannan Srinivasan, and Kaifu Zhang. Design of platform reputation systems: Optimal information disclosure. *Marketing Science*, 42(3):500–520, 2023.

[119] Jun-ji Shih, Chao-cheng Mai, and Jung-chao Liu. A general analysis of the output effect under third-degree price discrimination. *The Economic Journal*, 98(389):149–158, 1988.

[120] Naman Shukla, Arinbjörn Kolbeinsson, Ken Otwell, Lavanya Marla, and Kartik Yellepeddi. Dynamic pricing for airline ancillaries with customer context. In *Proceedings of the 25th ACM SIGKDD International Conference on knowledge discovery & data mining*, pages 2174–2182, 2019.

[121] Andy Smith. How the right incentive can help you generate more reviews. `https://www.powerreviews.com/what-incentives-generate-reviews/`, Nov 2022.

[122] Rakuten Insight Statista. Major reasons for stopping mobile dating app usage among respondents in hong kong as of february 2022. Feb 2022.

[123] Xuanming Su. Intertemporal pricing with strategic customer behavior. *Management Science*, 53(5):726–741, 2007.

[124] Arun Sundararajan. Nonlinear pricing of information goods. *Management Science*, 50(12):1660–1673, 2004.

[125] Steven Tadelis. Reputation and feedback systems in online platform markets. *Annual Review of Economics*, 8:321–340, 2016.

[126] NPR. Dating app burnout: When swiping becomes a chore. June 2019.

[127] Pew Research Center. Lesbian, gay and bisexual online daters report positive experiences – but also harassment. April 2020.

[128] Pew Research Center. The virtues and downsides of online dating. Feb 2020.

[129] Power Reviews. Beyond the basics: Key information shoppers look for in product reviews. Feb 2020.

[130] R Guy Thomas. Non-risk price discrimination in insurance: market outcomes and public policy. *The Geneva Papers on Risk and Insurance-Issues and Practice*, 37(1):27–46, 2012.

[131] Hal Varian. Price discrimination. In R. Schmalensee and R. Willig, editors, *Handbook of Industrial Organization*, volume 1, chapter 10, pages 597–654. Elsevier, 1 edition, 1989.

[132] Hal R Varian. Price discrimination and social welfare. *The American Economic Review*, 75(4):870–875, 1985.

[133] Hal R. Varian. Buying, sharing and renting information goods. *The Journal of Industrial Economics*, 48(4):473–488, 2000.

[134] Yue Wu, , and V Padmanabhan. The strategy puzzle of subscription based dating sites. *Harvard Business Review*, Jan 2019.

[135] Yue Wu, Kaifu Zhang, and V. Padmanabhan. Matchmaker competition and technology provision. *Journal of Marketing Research*, 55(3):396–413, 2018.

[136] Yingkang Xie and Huaiyu Zhu. Platform leakage: Incentive conflicts in two-sided markets. Technical report, Working Paper, https://questromworld. bu. edu/platformstrategy/wp-content ..., 2023.

[137] Zibin Xu and Anthony J Dukes. Price discrimination in a market with uninformed consumer preferences. *Available at SSRN: https://ssrn.com/abstract=2777081*, 2016.

[138] Yelp. Yelp - Company - Fast Facts. `https://www.yelp-press.com/company/fast-facts/default.aspx`, 2024. [Accessed 03-06-2024].

[139] Luisa Zhou. Online Review Statistics: The Definitive List (2024 Data). `https://www.luisazhou.com/blog/online-review-statistics/`, 2024. [Accessed 05-06-2024].

[140] F Zhu and M Iansiti. Why some platforms thrive and others don't. *Harvard Business Review*, 2019.

[141] Serhan Ziya, Hayriye Ayhan, and Robert D Foley. Relationships among three assumptions in revenue management. *Operations Research*, 52(5):804–809, 2004.

[142] Phillip Zukin. Yelp adds numerical ratings to help consumers easily compare multiple businesses. Technical report, 2023.