STRATEGIC CHALLENGES IN THE DIGITAL ADVERTISING ECOSYSTEM

by

Manmohan Aseri

APPROVED BY SUPERVISORY COMMITTEE:

____

Vijay S. Mookerjee, Co-Chair

____

Milind Dawande, Co-Chair

____

Ganesh Janakiraman

____

Harpreet Singh

Dedicated to my wife, daughter, and parents.

STRATEGIC CHALLENGES IN THE DIGITAL ADVERTISING ECOSYSTEM

by

MANMOHAN ASERI, BTech

DISSERTATION

Presented to the Faculty of

The University of Texas at Dallas

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY IN

MANAGEMENT SCIENCE

THE UNIVERSITY OF TEXAS AT DALLAS

August 2018

ACKNOWLEDGMENTS

First of all, I would like to thank my advisors, Professors Vijay Mookerjee, Milind Dawande, and Ganesh Janakiraman, for their training and guidance. I am extremely thankful to Professor Harpreet Singh for his constant help and support, and for agreeing to be part of my dissertation committee. I am grateful to Professors Atanu Lahiri, Nanda Kumar, and Upender Subramanian for their help during my job search. I am thankful to Professor Varghese Jacob for his guidance and also for his seminar on Information Systems. The learning from the seminar has greatly influenced my thinking and research, for which I am also indebted to the participants of the seminar.

I am extremely thankful to my wife, Ankita, for bearing with me during the PhD program, and giving me her unconditional support. I thank my parents, Mr. Laxminarayan Aseri and Mrs. Parvati Aseri, for encouraging me to pursue the path of academics. A special thanks to my brother and sisters for taking care of our parents back in India; this has definitely relieved me and allowed me to focus on my research here.

I thank my colleagues Harish Guda and Sameer Mehta for providing me with invaluable help during my job search. I also thank my friends Varun Gupta and Shivam Gupta for introducing me to Management Science and for encouraging me to apply to the doctoral program at UTD. I would also like to thank my office mates, Cheng Nie, Chenzhang Bao, and Danish Saifee, for all the fun and cooperation over the years. I also thank Jay Samuel for all those parties to keep the fun alive in the PhD program.

May 2018

STRATEGIC CHALLENGES IN THE DIGITAL ADVERTISING ECOSYSTEM

Manmohan Aseri, PhD
The University of Texas at Dallas, 2018

Supervising Professors: Vijay S. Mookerjee, Co-Chair
Milind Dawande, Co-Chair

This dissertation addresses three issues faced by different stakeholders in the digital advertising ecosystem. In the first chapter, we analyze a problem faced by Mobile-Promotion platforms, which enable advertisers (individual users or businesses) to directly launch their personalized mobile advertising campaigns. These platforms contract with advertisers to provide a certain number of impressions on mobile apps in their desired sets of geographic locations (usually cities or zip codes) within their desired time durations (for example, a month); the execution of each such contract is referred to as a campaign. To fulfill the demands of the campaigns, the platform bids in real-time at an ad exchange to win mobile impressions arising over the desired sets of locations of the campaigns and then allocates the acquired impressions among the ongoing campaigns. The core features that characterize this procurement problem – supply is uncertain, supply cannot be inventoried, and there are demand-side commitments to be met – are applicable to a variety of other business settings as well. Our analysis in this paper offers near-optimal policies for a static model representing a subscription-based setting, where customers provide information of their campaigns in advance to the platform. In the second chapter, we generalize our analysis of the first chapter and consider a dynamic model of campaign arrivals. The dynamic model represents a setting where campaigns arrive randomly and the platform reacts to these arrivals in real time; for this model, our rolling-horizon policy periodically re-computes the platform's procurement (or bidding) and allocation decisions. We establish performance bounds on our policies for both models and demonstrate the attractiveness of these bounds on real data. By isolating important structural features of a given set of campaigns, we discuss practical implementation issues and offer procurement-policy recommendations for a variety of settings based on these features. In the third chapter, we consider an ongoing issue of ad-blocking and analyze its impact on websites and consumers. We also propose strategies and insights that websites can use to react to ad-block users. We show that the website can increase its revenue by discriminating between regular and ad-block users via the ad-intensities shown to them. More interestingly, we find that the discriminatory power bestowed on the website by ad-blockers can also benefit its users when their outside option is not very attractive. Thus, the advent of ad-blockers can lead to a win-win for both the website and its users. Finally, we propose a superior *selective-gating* strategy in which only a fraction of ad-block users are gated. We establish the robustness of our conclusions under several enhancements to our base setting: (a) heterogenous profitabilities from regular users and ad-block users, (b) endogenous adoption of ad-blockers, (c) the presence of a subscription option, and (d) negative externality due to increased traffic. Our analysis ends with recommendations for three stakeholders in this problem, namely, publishers, web-browser developers, and policy makers.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1

INTRODUCTION

The ubiquity of online media such as smart phones, tablets, and personal computers, has led to their increasing use for advertising over the past decade (Central Market Research, 2012; Lieberman, 2013). In particular, smart phones and tablets are rarely shared by multiple individuals and can be tracked, thereby offering tremendous opportunities for targeted advertising. Not surprisingly, mobile advertising is leading the rise in media ad spending, with annual revenues in the US of about $18 billion in 2014 and projected to exceed $50 billion by 2018 (eMarketer, 2014).

The focus of the current study is on advertising on a mobile device (e.g., a smart phone, or a tablet). Specifically, we consider ads that are displayed on a mobile application (hereafter abbreviated as an *app*), such as an app for weather, stocks, or a game. This form of advertising is on the increase. Here, *advertisers* – who provide ads created to promote products or services, and ultimately generate demand for advertising space – and *publishers*, i.e., owners of the mobile apps that supply the space for ad-display, often interact with each other through their respective agents. Advertisers are represented by *demand aggregators* (e.g., Cidewalk, Sitescout) who provide advertisers with access to a variety of publishers for the appropriate exposure of their ads. On the other hand, publishers are represented by *supply aggregators* (e.g., Tribal Fusion, Chitika) who help monetize the space owned by publishers and earn revenue for them.

The opportunity to display an ad on an app is referred to as an *impression*. The arrival (or generation) of impressions is determined by a host of factors, including population density, time of day, and weather. Till recently, most of the supply of impressions was sold on a contractual basis. That is, supply aggregators entered into relatively long-term contractual arrangements with app owners to sell their space for the display of ads. These contracts were often drawn on a revenue-sharing basis, i.e., app owners got a proportion of the revenue generated from their ad-space. In recent times, with the advent of *mobile ad exchanges* (e.g., DoubleClick, Opera), the supply of impressions on apps is fast becoming a commodity. There are no long-term contracts; rather, each opportunity to display an ad is auctioned off on a mobile ad exchange.

The growth of mobile ad exchanges has also led to changes in the demand-side of the industry. Demand aggregators are now able to directly buy supply from exchanges rather than accessing the supply via a supply aggregator. While the advertising needs of large advertisers (such as GE or Sony) are often addressed in an "ad agency" mode providing end-to-end service, demand aggregators are moving to a "managed-service" mode for small advertisers. Here, advertisers directly launch their individual campaigns using a *mobile-promotion platform* (a demand aggregator), which accepts the operational responsibility of procuring impressions from the desired locations and at the desired times to fulfill these campaigns. This "managed-service" mode is suitable for thousands of small advertisers without deep pockets and offers an attractive advertising solution for a hitherto under-served segment of the market. Further, managed service can be of two types: *subscription-based* and *"walk-in"*. In the former, advertisers provide information of their campaigns (when they arrive, how many impressions they need, etc.) in advance to the platform. In the latter, the advertisers start their campaigns dynamically and the platform reacts in real time to satisfy the requirements. In Chapter 2, we analyze the optimal bidding policies for a platform that offers a subscription-based service. Then, in Chapter 3, we analyze the bidding policies for a platform that offers walk-in service.

A mobile-promotion platform can be seen as an online substitute for the traditional (offline) Yellow Pages. As with Yellow Pages, such platforms are ideal for small local businesses. However, medium and large firms also use them to run advertising campaigns targeted in specific geographic locations. A prominent example of a mobile-promotion platform is Cidewalk, the company we interacted with in defining the problems addressed in this paper. Other major mobile-promotion platforms include Sitescout (`http://www.sitescout.com/`) and ExactDrive (`http://www.exactdrive.com/`).

Mobile-Promotion platforms scale well and provide access to supply via the use of mobile ad exchanges. The platform's key expertise is in the ability to bid intelligently to procure (i.e., win) impressions on a mobile ad exchange. Also, there needs to be a high level of integration between the ad exchange and the platform to complete the bidding process in real-time (Real-Time Bidding or RTB) and render the winner's ad on an app. Such technical expertise is usually not possessed by small advertisers. Hence there is a niche for firms that possess sophisticated integration skills and fast, real-time analytic abilities to buy supply at affordable prices and deliver ad campaigns at a net profit.

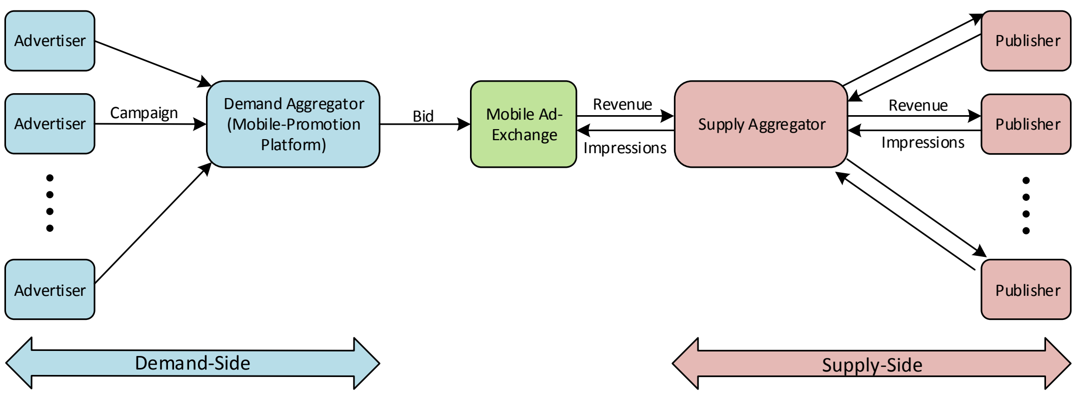

Figure 1.1. *The Mobile-Promotion Supply Chain*: Advertisers start their campaigns on a mobile-promotion platform, which executes these campaigns by procuring the required impressions on a mobile ad-exchange, where publishers come to monetize their supply of space for ad-display.

Figure 1.1 depicts the operational details of the mobile-promotion supply chain. The mobile-promotion platform enters into contracts with advertisers (individual users or businesses) under which it has to provide a certain number of impressions on mobile apps in their desired sets of geographic locations (usually a set of cities or zipcodes) within their desired time durations (e.g., a week or a month). We will refer to the execution of a contract with an individual advertiser as a *campaign*. Thus, the input corresponding to a campaign includes a target number of impressions, a *set* of locations in which the impressions have to be provided, and a time duration of execution. There are multiple such campaigns being executed simultaneously. The sets of locations of the campaigns as well as their time durations may intersect. The platform bids on ad-exchanges to win the impressions needed (to place ads) to fulfill these campaigns: The more the bid on an impression, the higher is the probability of winning that impression. The objective of the platform is to win enough impressions and allocate these impressions among the ongoing campaigns to fulfill the requirements of the campaigns at minimum cost. Our goal in this paper is to address the problem faced by the platform of optimizing its procurement (or bidding) and allocation policy for mobile ad delivery to support its multiple ad campaigns.

In Chapter 4, we address the problem of ad-blocking, faced by websites and analyze the impact of ad-blocking. Online advertisements have been pivotal in keeping the internet largely free of cost for users. Many websites – small and large – rely solely on advertising revenue for their survival. While it is a lifeline for a free internet, online advertising has been widely criticized for deteriorating consumer experience by displaying an excessive amount of irrelevant ads to users (Wall Street Journal, 2014). Perhaps as an inevitable systemic retaliation, recent advancements in technology have allowed consumers of digital content to block advertisements by using a type of software that is now commonly referred to as an *Ad-Blocker* (e.g., Adblock Plus, see Figure 1.2). A typical ad-blocker works by blocking ads at their source, i.e., by preventing communication to ad-servers that deliver ads. Therefore, the use of ad-blockers not only improves the visual experience of the user, but also results in other advantages, including reduced data download and website loading time, and efficient use of battery power. It has been reported that websites load nearly four times faster with an ad blocker and use about 50% less data (Wall Street Journal, 2015).

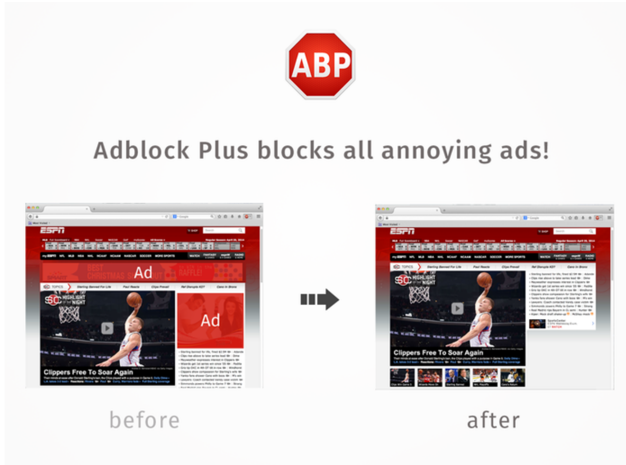

Figure 1.2. The ad-cleansing action of an ad-blocker (source: Adblock Plus).

Motivated by these benefits, companies like Apple and Samsung – that do not rely heavily on ad revenues – have welcomed ad-blockers on their mobile devices by adding the functionality to integrate third-party ad-blockers. According to a highly cited recent report by PageFair,[1] ad-blocking penetration in the U.S. is currently around 18% and in some other countries, such as Indonesia, has reached above 50%. It is estimated that publishers worldwide lost nearly $41.4 billion in 2016 to the use of ad-blockers.

Most ad-blockers come with a "white-listing" feature, that allows web-users to create a list of websites, referred to as a *white-list*, in which the ad-blocker will not prevent ads from appearing. If an ad-blocker is detected on a user's device, the target website can ask the user to white-list the website, failing which access to the website can be denied; examples include `http://www.forbes.com` and `http://www.businessinsider.com`. Thus, websites do have a choice in allowing or denying access to ad-block users. This choice raises a genuine predicament: on the one hand, ad-block users do not generate revenue for the website; on the other hand, denying them access can shrink the user base and adversely affect the popularity of the website, ultimately reducing traffic over the long run

[1] https://pagefair.com/blog/2017/adblockreport/

(we refer to this phenomenon as a *network effect*). Indeed, many publishers who reacted to the increase in ad-block users by denying them access quickly stepped back and stopped doing so (e.g., `http://www.ratemyprofessors.com`). To our knowledge, no clear strategy has yet emerged to resolve this dilemma.

We model the decision problem for a publisher facing two user segments: regular users (hereafter, referred to as *regulars*) and *ad-blockers*.[2] Regulars are potential users of the website who do not use any ad-block software, whereas ad-blockers are potential users of the website who also use some ad-block software. The first-level decision or *gating strategy* is whether to (a) allow ad-free access to ad-block users or (b) require white-listing by ad-block users in order to allow them access. In (a), the second-level decision for the website is the level of advertising (*ad intensity*) to regular users. In (b), the second-level decisions are the ad-intensities to be used for regular users and for ad-block users.[3] The website's revenue depends directly on the ad-intensity(ies) it chooses and the traffic it generates from each of the two user segments. Moreover, the traffic it generates depends on the *net utility* that each individual user receives from visiting the website and the value of the best alternative to the website (i.e., the outside option). The net utility offered by the website depends on (i) the intrinsic value of the website's content, (ii) the value users obtain due to network effects, i.e., value derived as a result of the amount of traffic/popularity of the website, and (iii) a negative utility or cost incurred due to the presence of ads, driven by the overall past experience of ads as a nuisance; regular and ad-block users differ in terms of these *ad-viewing costs*. Using this model, we derive an optimal gating and ad-intensity strategy for the website. We also solve an identical model for a world without ad-block software. We show analytically that the presence of such

[2]When no confusion arises in doing so, we use the term "ad-blocker" to interchangeably refer to ad block software or a user who has installed such software.

[3]Websites can detect and track white-listers; see, e.g., (DVorkin, 2016).

software can be used by the website to increase its revenue by discriminating between regular and ad-block users in terms of the ad-intensities shown to them. Clearly, this was not possible in the absence of such software.

One naturally wonders if the website's benefit in the post-ad-block world comes at the expense of consumer welfare. Interestingly, we establish the precise condition – the utility users get from their outside option is lower than a threshold – under which *both* the website and its consumers benefit after the advent of ad-block software. Intuitively, when the website is able to customize ad-intensities for regulars and ad-blockers, the total traffic to the website increases. When the outside option offers a low utility, these additional users not only obtain a higher net utility (relative to the pre-ad-block world) but also increase the value of the website, which in turn also increases the net utility of the existing users.

Instead of gating either all or none of the users (*integral gating*), one can consider a more-general *selective gating* strategy in which only a fraction of the ad-block users are asked to white-list the website. We show that this can increase the website's revenue significantly from that under integral gating. This analysis also helps us derive insights on the optimal selective gating strategy, with respect to the strength of the network effect and the *ad-blocking rate* (the fraction of the potential user-population that uses ad-blockers).

CHAPTER 2

PROCUREMENT POLICIES FOR MOBILE-PROMOTION PLATFORMS[1]

Unlike classical settings in Operations Management where uncertainty in demand is addressed via supply-side commitments, a mobile-promotion platform faces a situation where it has made firm delivery commitments to customers but its supply of impressions is uncertain. There are two important distinctions here from the traditional literature on delivery commitments under supply/demand uncertainty (see, e.g., Chapters 8 and 10 in Tayur et al. 2012) and that on inventory management under service-level constraints (e.g., Bookbinder and Tan 1988; Bashyam and Fu 1998): First, the platform's uncertain supply of impressions cannot be inventoried. In other words, the platform cannot win impressions and store them to satisfy requirements in the future. Second, supply uncertainty arises here due to reasons unlike those in the classical OM literature, e.g., random yields (Yano and Lee 1995; Wang and Gerchak 1996): the uncertain arrival of impressions followed by real-time competitive bidding to win these impressions. Another important distinguishing feature is that the platform faces a multi-campaign problem over multiple locations: On the one hand, impressions won from a location can be distributed to multiple campaigns; this is akin to distribution systems, where a manufacturer allocates supply to multiple retailers. On the other hand, a single campaign's demand can be satisfied by sourcing impressions from multiple locations – this is similar to having multiple suppliers for a product. Further, since the sets of locations differ across campaigns, the situation becomes similar to one where customers (campaigns) desire subsets of partially-substitutable items (impressions for a campaign are sourced only from a specific subset of

[1]Republished with permission of INFORMS, from "Procurement Policies for Mobile-Promotion Platforms," M. Aseri, M. Dawande, G. Janakiraman, and V. Mookerjee, forthcoming in Management Science, October 2017; permission conveyed through Copyright Clearance Center, Inc.

locations, but they could come from any location in the subset). Another related stream from the supply chain literature is that on managing inventories in the face of uncertain demand and a spot market for procurement (Haksöz and Seshadri 2007; Chen et al. 2013). One feature of our problem which further differentiates our work from the traditional Operations Management literature is that the network structure of our problem (see Figure 1.1) changes much more frequently (whenever a new campaign arrives or leaves) than in traditional settings.

As far as the application context is concerned, our paper is naturally related to the growing literature on targeted advertising and, more generally, display advertising. Mobile-promotion platforms help advertisers exploit both the location of their audience and the time when their advertisements are sent. The efficacy of both these factors has been highlighted in several studies; see e.g., Forman et al. (2009), Luo et al. (2013), Andrews et al. (2015), Fang et al. (2015). In the vast literature on display advertising, there are many recent studies that address the challenges faced by the main stakeholders in the digital advertising ecosystem. For example, Balseiro et al. (2014) address the problem of a publisher who has contracted with advertisers to deliver a certain number of impressions, but still has a choice of selling the impressions on an ad-exchange. The goal of the paper is to provide a policy to decide which impressions should be assigned to contracted advertisers and which should be sold on the ad-exchange. This problem is fundamentally different from the one we address: First, the problem in Balseiro et al. (2014) is from the viewpoint of a publisher while the problem we address is from the perspective of a platform on which several advertisers execute their campaigns and the goal of the platform is to satisfy these campaigns (i.e., deliver the required number of impressions) at minimum cost. Second, Balseiro et al. (2014) assume that supply is plentiful, i.e., a sufficient number of impressions are guaranteed to arrive to satisfy the requirement of the contracted advertisers. On the other hand, the very core of the problem addressed in our paper is that the supply (the arrival of impressions) and the procurement of impressions (winning them) are both uncertain. We mention three other papers: A problem of pricing impressions faced by a publisher, who charges based on a cost-per-click pricing scheme, is addressed in Najafi-Asadolahi and Fridgeirsdottir (2014). Turner (2012) addresses an ad-network's problem of allocating impressions generated by multiple demographic groups of users to multiple advertisers, with the objectives of offering viewers a variety of ads and making the execution highly predictable. Chickering and Heckerman (2003) address a problem faced by a publisher of maximizing the click-through-rate, subject to constraints that guarantee a minimum volume of impressions for each advertiser.

The Computer Science literature investigates several other challenges encountered by the different stakeholders in display advertising. For instance, (i) the design of contracts between publishers and advertisers (see, e.g., Babaioff et al. 2009, Constantin et al. 2009, and Feige et al. 2009), (ii) ad serving strategies for publishers who contract with multiple advertisers (see, e.g., Mahdian et al. 2007, Feldman et al. 2009, Kesselheim et al. 2014, and Esfandiari et al. 2015), (iii) the tradeoff faced by publishers between selling impressions at ad exchanges versus selling them via long-term contracts with advertisers (see, e.g., Ghosh et al., 2009, Balseiro et al. 2014, Arnosti et al. 2016), (iv) revenue-maximizing strategies for ad exchanges when there is a significant gap between the highest and second-highest bids for certain impressions (see, e.g., Mirrokni et al. 2010, Celis et al. 2014), (v) learning of advertisers' valuations for impressions through repeated interactions between an ad exchange and advertisers (see, e.g., Amin et al. 2013, Kanoria and Nazerzadeh 2014). For a broader discussion of these and other challenges, we refer the reader to the recent survey by Korula et al. (2016).

Finally, the analysis of our rolling-horizon policy for the platform's procurement and allocation decisions makes a useful contribution to the existing literature on re-optimization methods. We mention three recent studies: Jasin and Kumar (2012) consider a network revenue management problem with customer choice and exogenous prices, and develop a performance guarantee for a class of heuristics that periodically re-solve deterministic linear programs. (Jasin, 2014) addresses a seller's goal of maximizing total expected revenue for a dynamic pricing problem with finite inventories, finite selling horizon, and stochastic demands, and establishes performance guarantees for heuristics that require few or no re-optimizations. Motivated by applications in display advertising and network revenue management, Ciocan and Farias (2012) consider a class of dynamic allocation problems with unknown and volatile demand. The authors establish performance guarantees for a simple algorithm that uses a combination of re-optimization and "robust" forecasting.

2.1 Analysis of a Static Model of Campaign Arrivals

In this section, we describe and study a model in which information about all the campaigns, i.e., their time durations, demands, and locations of interest, is available at the beginning of the planning horizon. Thus, this model is a *static model* of campaign arrivals. The formulation and analysis of this model will serve as a useful stepping stone for the *dynamic model* studied in Chapter 3.

Let $\mathcal{L} = \{1, 2, \ldots, L\}$ denote the set of locations[2] (say, zip-codes) in the geographic region considered. Each campaign is interested in displaying its ad over a specific set of locations on mobile devices when users open apps – we refer to the event of a mobile device user, in one of the desired locations, opening an app as an *impression*. Thus, there are $2^L - 1$ campaign-types based on the set of locations a campaign is interested in. Let $\mathcal{C} = \{1, 2, \ldots, C\}$ denote this set, where $C = 2^L - 1$. For every $c \in \mathcal{C}$, we use $\mathcal{L}_c$ to denote the set of locations of interest to campaign-type c. For ease of exposition, we assume

[2]The notion of a location here can also be used to model the finer attributes of an impression such as the (location, app) combination and/or the demographic segment of the mobile user.

that one campaign of each type arrives at the beginning of the horizon;[3] thus, we use "campaign c" to refer to the type-c campaign which arrives. For every $c \in \mathcal{C}$, let M_c denote the number of impressions required by campaign c.

There are three units of time which we use in our model. The planning horizon consists of a set of S periods denoted by $\mathcal{S} = \{1, 2, \ldots, S\}$, where a time period is typically a day. Each time period consists of I time blocks indexed by $i = 1, 2, \ldots, I$, where each block is typically a segment of a day, say early morning, mid-morning, etc.. Let $\mathcal{I} = \{1, 2, \ldots, I\}$ be the set of distinct time blocks. Each block consists of many time slots, where a time slot is typically a millisecond. For notational convenience, each block is assumed to consist of Δ time slots. Then, the planning horizon consists of the set of time slots $\mathcal{T} := \{1, 2, \ldots, SI\Delta\}$. We also use the notation $\mathcal{T}_s$ to denote the set of time slots in period s and $\mathcal{T}_{s,i}$ to denote the set of time slots in block i of period s.

The length of a time slot is chosen to be sufficiently small to ensure that at most one impression arrives in one time slot. For any time slot $t \in \mathcal{T}_{s,i}$, let $q_{i,l}$ be the probability that an impression from location l arrives in that time slot. Impression arrivals are probabilistically independent across time slots. The impressions are auctioned in real-time on an ad-exchange. The platform, and also other advertisers or companies acting on behalf of advertisers, bid for these impressions. The highest bidder wins the impression and pays the price that she bids to the exchange. Clearly, the higher the bid, the greater the probability of winning an impression. We model this using a *win-curve*, that is, a function $p_{i,l}(b) : [0, b_{i,l}^{max}] \to [0, 1],\ l \in \mathcal{L},\ i \in \mathcal{I}$, which means the following: Given that an impression arrives from location l in some time slot $t \in \mathcal{T}_{s,i}$ (for some $s \in \mathcal{S}$ and $i \in \mathcal{I}$), $p_{i,l}(b)$ specifies the probability of winning that impression by bidding an amount b. It is

[3]The assumption that each campaign arrives at the beginning of the planning horizon is purely for expositional convenience. The analysis in this section easily extends to the situation when there are (possibly) multiple campaigns of each type and they start at different times.

reasonable and convenient to assume that $b_{i,l}^{max}$ is a large enough value that an impression will definitely be won with a bid of that value.

For every $c \in \mathcal{C}$, let K_c denote the number of periods over which campaign c lasts; that is, this campaign's duration is K_c periods. Thus, $t_c := K_c I \Delta$ is the last time slot in campaign c's duration. That is, campaign c requires a delivery of M_c impressions from the set of locations, $\mathcal{L}_c$, within the first K_c periods (i.e., first t_c time slots). For every campaign, the platform commits to meeting this requirement with a probability α, whose value is close to 1.[4]

The platform's decisions are as follows: (i) *Procurement Policy* – The bids to place through time on the impressions that (possibly) arrive from the desired locations. (ii) *Allocation* – If an impression is won from a certain location and if several campaigns are interested in that location, then the platform also chooses the campaign to which that impression is to be assigned.

Notice that the platform's bids can be *dynamic*; for example, if many impressions have been won early on within the campaign time-duration, then the platform may start bidding low on subsequent impressions. The platform's objective is to minimize its expected cost subject to the constraint that it obtains at least M_c impressions with a probability of α or more, for each campaign $c \in \mathcal{C}$. We denote this problem by $\mathcal{P}_{static}(\vec{M}, \alpha)$, where $\vec{M} \in \mathbb{R}_+^C$ is the vector with elements $M_c, c \in \mathcal{C}$.

For our mathematical formulation of $\mathcal{P}_{static}(\vec{M}, \alpha)$, it is useful to define the *bid curve* $b_{i,l}(x) : [0,1] \to [0, b_{i,l}^{max}]$, $i \in \mathcal{I}$, $l \in \mathcal{L}$, as the inverse of the function $p_{i,l}(\cdot)$; that is, $b_{i,l}(x) = p_{i,l}^{-1}(x)$. Thus, $b_{i,l}(x)$ is the bid required during time-block i at location l, to ensure a winning probability of x. We also define $f_{i,l}(x) = x b_{i,l}(x)$; this is the expected cost associated with choosing a win-probability of x for an impression at location l in time-block i.

[4]In practice, it is not possible for the platform to choose $\alpha = 1$, because there is always a chance (albeit small) that very few impressions arrive during the campaign duration (e.g., a natural calamity may cause this). If the platform fails to meet the demand for a campaign, then the payment is refunded to the advertiser; however, with the choice of a high value of α, this event is quite rare.

It should be clear that our problem can also be formulated as one in which the platform's decisions are the win-probabilities to use on the impressions that arrive (instead of the corresponding bid amounts). We now proceed to present the mathematical formulation of $\mathcal{P}_{static}(\vec{M}, \alpha)$.

Consider any period s in the planning horizon $\mathcal{S}$ and any time block $i \in \mathcal{I}$. For every $t \in \mathcal{T}_{s,i}$, let $j_{t,c}$ denote the number of impressions won by campaign c before the start of time-slot t. Thus, $\vec{j}_t = (j_{t,c} : c \in \mathcal{C})$ denotes the state vector at the beginning of time-slot t; this vector captures the information on the number of impressions won by the various campaigns. Let $\mathcal{J}_t$ denote the set of all feasible state vectors in time-slot t.

A policy π is defined by a set of target win-probability functions $\{x^{\pi}_{t,l} : \mathcal{J}_t \to [0,1],\ t \in \mathcal{T},\ l \in \mathcal{L}\}$ and a set of allocation functions $\{y^{\pi}_{t,l,c} : \mathcal{J}_t \to [0,1],\ t \in \mathcal{T},\ l \in \mathcal{L},\ c \in \mathcal{C}\}$. If an impression arises from location l in time-slot t, then $x^{\pi}_{t,l}(\vec{j}_t)$ denotes the win-probability targeted by π for this impression and $y^{\pi}_{t,l,c}(\vec{j}_t)$ denotes the probability with which policy π wins this impression and allocates it to campaign c.[5] We also assume an upper bound $x_{max} \in (0, 1]$ on the target win-probabilities $\{x^{\pi}_{t,l}\}$. In practice, most of the campaigns come from metropolitan areas, where the supply of mobile impressions is plentiful. Therefore, the platform typically needs to procure only a small fraction of the total supply of impressions. For instance, Cidewalk does not need to target a win-probability higher than 0.05 from any location in any time slot (Cidewalk, Inc., 2016). From a theoretical perspective, we will see that the provable cost performance of the policy we propose depends on x_{max}.

Next, we explain how we model the random outcomes associated with the arrival, the winning, and the allocation of impressions. Consider any period $s \in \mathcal{S}$ and block $i \in \mathcal{I}$. For every slot $t \in \mathcal{T}_{s,i}$, let U^{imp}_t be a random variable that takes any value $l \in \mathcal{L}$ with

[5] We note that the decisions $x^{\pi}_{t,l}$ and $y^{\pi}_{t,l,c}$ depend on the history until time slot t only through $\vec{j}_t$. This is because of the assumed independence of the arrival of impressions across time slots (over all time periods).

probability $q_{i,l}$ (denoting the event that an impression arrives in location l in time slot t) and the value 0 with probability $1 - \sum_{l \in \mathcal{L}} q_{i,l}$ (denoting the event that no impression arrives in time slot t). We assume that these random variables are independent across time slots. To model the winning and the allocation of impressions, we use a set of $U[0,1]$ random variables (this is a standard and convenient mathematical device) as follows.

Let $\{U_t^{win} : t \in \mathcal{T}\}$ and $\{U_t^{alloc} : t \in \mathcal{T}\}$ denote two sets of mutually independent i.i.d $U[0,1]$ random variables, independent of $\{U_t^{imp} : t \in \mathcal{T}\}$. If an impression arises from location l in time-slot t, we say that this impression is won by policy π if and only if $U_t^{win} \leq x_{t,l}^{\pi}(\vec{j}_t)$. If this impression is won, then it is allocated to campaign c if U_t^{alloc} is between $\sum_{\hat{c}=1}^{c-1} y_{t,l,\hat{c}}^{\pi}(\vec{j}_t)/x_{t,l}^{\pi}(\vec{j}_t)$ and $\sum_{\hat{c}=1}^{c} y_{t,l,\hat{c}}^{\pi}(\vec{j}_t)/x_{t,l}^{\pi}(\vec{j}_t)$. This ensures that, in time slot t, an impression is won for campaign c from location $l \in \mathcal{L}_c$ with probability $y_{t,l,c}^{\pi}(\vec{j}_t)$. Let $\vec{U}_t = \left(U_1^{imp}, U_1^{win}, U_1^{alloc}, U_2^{imp}, U_2^{win}, U_2^{alloc}, \ldots, U_t^{imp}, U_t^{win}, U_t^{alloc}\right)$ be the history of these random variables until time-slot t (including t) and let $\mathcal{U}_t$ be the set of all possible vectors $\vec{U}_t$. Also, let $j_{t,c}^{\pi}(\vec{u}_{t-1})$ represent the number of impressions procured by policy π for campaign c by the start of time slot t when the realized history $\vec{U}_{t-1} = \vec{u}_{t-1}$. Let $\vec{j}_t^{\pi}(\vec{u}_{t-1})$ be the vector containing elements $j_{t,c}^{\pi}(\vec{u}_{t-1}),\ \forall c \in \mathcal{C}$.

From the definition of $x_{t,l}^{\pi}(\vec{j}_t)$ and $y_{t,l,c}^{\pi}(\vec{j}_t)$, we note that $x_{t,l}^{\pi}(\vec{j}_t) = \sum_{c \in \mathcal{C}} y_{t,l,c}^{\pi}(\vec{j}_t),\ \forall\, \vec{j}_t \in \mathcal{J}_t$. Furthermore, since campaign c is only interested in the set of locations $\mathcal{L}_c$, we have $y_{t,l,c}^{\pi}(\vec{j}_t) = 0,$ if $l \notin \mathcal{L}_c,\ \forall c \in \mathcal{C},\ \vec{j}_t \in \mathcal{J}_t$. For any policy π, let

$$f^{\pi} := \sum_{s \in \mathcal{S}} \sum_{i \in \mathcal{I}} \sum_{t \in \mathcal{T}_{s,i}} \sum_{l \in \mathcal{L}} \mathbb{E}_{\vec{U}_{t-1}} \left[q_{i,l} f_{i,l}\left(x_{t,l}^{\pi}\left(\vec{j}_t^{\pi}(\vec{U}_{t-1})\right)\right)\right]$$

denote the expected cost of policy π.

The platform's problem $\mathcal{P}_{static}(\vec{M}, \alpha)$ can now be written as follows:

$$
\mathcal{P}_{static}(\vec{M},\alpha):\begin{cases}
\min_\pi \; f^\pi & (2.1)\\
\text{subject to: } \sum_{c\in\mathcal{C}} y^\pi_{t,l,c}(\vec{j}) = x^\pi_{t,l}(\vec{j}),\ \forall\, t\in\ \mathcal{T},\ \vec{j}\in J_t,\ l\in\mathcal{L}, & (2.2)\\
j^\pi_{t+1,c}(\vec{u}_t) \;=\; j^\pi_{t,c}(\vec{u}_{t-1}) \;+\; \sum_{l\in\mathcal{L}_c} \mathbb{1}(U^{imp}_t = l) & (2.3)\\
\times\mathbb{1}\left(U^{win}_t \le x^\pi_{t,l}\left(\vec{j}^\pi_t(\vec{u}_{t-1})\right)\right) & \\
\times\mathbb{1}\left(\frac{\sum_{\hat{c}=1}^{c-1} y^\pi_{t,l,\hat{c}}\left(\vec{j}^\pi_t(\vec{u}_{t-1})\right)}{x^\pi_{t,l}\left(\vec{j}^\pi_t(\vec{u}_{t-1})\right)} < U^{alloc}_t \le \frac{\sum_{\hat{c}=1}^{c} y^\pi_{t,l,\hat{c}}\left(\vec{j}^\pi_t(\vec{u}_{t-1})\right)}{x^\pi_{t,l}\left(\vec{j}^\pi_t(\vec{u}_{t-1})\right)}\right), & \\
\forall\, t\in\ \mathcal{T},\ c\in\mathcal{C},\ \vec{u}_t\in\mathcal{U}_t, & \\
x^\pi_{t,l}(\vec{j}),\ y^\pi_{t,l,c}(\vec{j})\in[0,x_{max}],\ \ \forall\, t\in\ \mathcal{T},\ \vec{j}\in\mathcal{J}_t,\ l\in\mathcal{L},\ c\in\mathcal{C}, & (2.4)\\
\mathbb{P}\left[j^\pi_{t_c+1,c}(\vec{U}_{t_c})\ge M_c\right] \;\ge\; \alpha,\ \forall\, c\in\mathcal{C}. & (2.5)
\end{cases}
$$

Equation (2.3) defines the transition of $j^\pi_{t,c}(\vec{u}_{t-1})$ to $j^\pi_{t+1,c}(\vec{u}_t)$. To understand this equation, note that the first indicator variable on its right side takes the value 1 if an impression arrived at location l in time slot t, the second indicator takes the value 1 if this impression is won by policy π, and the third indicator takes the value 1 if the impression is allocated to campaign c.

Our plan to analyze $\mathcal{P}_{static}(\vec{M},\alpha)$ is as follows: In the next subsection, we will study a relaxation of problem $\mathcal{P}_{static}(\vec{M},\alpha)$ and obtain an *optimal* policy for this relaxation. Then, in Section 2.1.2, we will use that policy to obtain a *near-optimal* policy for problem $\mathcal{P}_{static}(\vec{M},\alpha)$.

Note: We have assumed that the arrivals of impressions are independent across time slots. However, our analysis can be extended for an arbitrary stochastic process of impression arrivals; see Remark 2 in Section 2.1.2 for a discussion.

2.1.1 A Relaxation and Its Optimal Solution

We start by defining a new problem, similar to problem $\mathcal{P}_{static}(\vec{M},\alpha)$, in which the constraint is to deliver a certain number of impressions, say $\beta_c \ge 0$, *in expectation* for each campaign $c\in\mathcal{C}$. Let $\vec{\beta}$ be the vector with elements β_c, $c\in\mathcal{C}$. Formally, we denote by

$\mathcal{P}^{E}_{static}(\vec{\beta})$ the following problem:

$$\mathcal{P}^{E}_{static}(\vec{\beta}) : \begin{cases} \min_{\pi} \ f^{\pi} \\ \text{subject to: (2.2), (2.3), (2.4) and} \\ \mathbb{E}\big[j^{\pi}_{t_c+1,c}(\vec{U}_{t_c})\big] \ \geq \ \beta_c, \ \forall\, c \in \mathcal{C}. \end{cases} \tag{2.6}$$

Next, using Markov's inequality we observe that, for each $c \in \mathcal{C}$, the probabilistic guarantee in problem $\mathcal{P}_{static}(\vec{M}, \alpha)$, that is, the constraint $\mathbb{P}\big[j^{\pi}_{t_c+1,c}(\vec{U}_{t_c}) \geq M_c\big] \geq \alpha$ implies the inequality $\mathbb{E}\big[j^{\pi}_{t_c+1,c}(\vec{U}_{t_c})\big] \geq \alpha M_c$. Thus, for the choice $\vec{\beta} = \alpha\vec{M}$, problem $\mathcal{P}^{E}_{static}(\alpha\vec{M})$ is a relaxation of problem $\mathcal{P}_{static}(\vec{M}, \alpha)$. We now proceed to solve $\mathcal{P}^{E}_{static}(\vec{\beta})$ for any $\vec{\beta}$.

We assume that for $l \in \mathcal{L}$ and $i \in \mathcal{I}$, the expected cost $f_{i,l}(x)$ of ensuring a win-probability of x at location l during time block i, is a strictly increasing and strictly convex function of x. This is a reasonable assumption; for instance, this result can be easily derived if we assume that the win curve $p_{i,l}(b)$ – which defines the probability of winning an impression at location l during time block i by bidding an amount b – is strictly increasing and concave in b (or, equivalently, that the bid curve $b_{i,l}(x)$ is strictly increasing and convex in x). This latter assumption has been verified on real data in prior research; see e.g., Zhang et al. (2014). In Chapter 3, we validate the strict convexity of $f_{i,l}(x)$ on real data. Accordingly, we will henceforth assume that $f_{i,l}(x)$ is strictly increasing and strictly convex in x. We now introduce a *deterministic* problem $\hat{\mathcal{P}}(\vec{\beta})$ and then establish its equivalence to $\mathcal{P}^{E}_{static}(\vec{\beta})$; that is, we show that an optimal policy for $\mathcal{P}^{E}_{static}(\vec{\beta})$ can be obtained from an optimal solution to $\hat{\mathcal{P}}(\vec{\beta})$. For any vector $\hat{x} = (\hat{x}_{s,i,l} : s \in \mathcal{S},\ i \in \mathcal{I},\ l \in \mathcal{L})$, let $\hat{f}(\hat{x}) = \sum_{s\in\mathcal{S}} \sum_{i\in\mathcal{I}} \sum_{l\in\mathcal{L}} \Delta\, q_{i,l}\ f_{i,l}(\hat{x}_{s,i,l})$.

Let $\hat{\mathcal{P}}(\vec{\beta})$ denote the following problem:

$$\hat{\mathcal{P}}(\vec{\beta}) : \begin{cases} \min_{\hat{x},\hat{y}} \ \hat{f}(\hat{x}) \\ \text{subject to: } \sum_{c\in\mathcal{C}} \hat{y}_{s,i,l,c} \ = \ \hat{x}_{s,i,l}, \qquad \forall\, s \in \mathcal{S}, i \in \mathcal{I}, l \in \mathcal{L}, & (2.7) \\ \sum_{s=1}^{K_c} \sum_{i\in\mathcal{I}} \sum_{l\in\mathcal{L}_c} \Delta\, q_{i,l}\hat{y}_{s,i,l,c} \ \geq \ \beta_c, \ \forall\, c \in \mathcal{C}, & (2.8) \\ \qquad \hat{x}_{s,i,l},\ \hat{y}_{s,i,l,c} \in [0, x_{max}], \qquad \forall\, s \in \mathcal{S}, i \in \mathcal{I}, l \in \mathcal{L}, c \in \mathcal{C}. \end{cases}$$

Theorem 1. *Let $\big(\hat{x}^*(\vec{\beta}), \hat{y}^*(\vec{\beta})\big)$ be an optimal solution to $\hat{\mathcal{P}}(\vec{\beta})$. Let $\pi(\vec{\beta})$ denote the following policy:*

$$\pi(\vec{\beta}) := \begin{cases} x_{t,l}^{\pi(\vec{\beta})}(\vec{j}) = \hat{x}_{s,i,l}^*(\vec{\beta}), \ \ \forall\, t \in \mathcal{T}_{s,i}, \ s \in \mathcal{S}, \ i \in \mathcal{I}, \ l \in \mathcal{L}, \ \vec{j} \in \mathcal{J}_t, \\ y_{t,l,c}^{\pi(\vec{\beta})}(\vec{j}) = \hat{y}_{s,i,l,c}^*(\vec{\beta}), \ \forall\, t \in \mathcal{T}_{s,i}, \ s \in \mathcal{S}, \ i \in \mathcal{I}, \ l \in \mathcal{L}, c \in \mathcal{C}, \ \vec{j} \in \mathcal{J}_t. \end{cases}$$

Then, $\pi(\vec{\beta})$ is an optimal policy for $\mathcal{P}^E_{static}(\vec{\beta})$.

The proofs of all the technical results are in the Appendix. Two important observations about the policy $\pi(\vec{\beta})$ in Theorem 1 are relevant here: (i) This policy is a *static* (i.e., state-independent) policy. That is, the platform's target win-probability at a given time does not depend upon the state of the system at that time, namely the number of impressions won for the individual campaigns. This is an appealing feature, since the platform does not need to alter its bidding intensity either with time (within a time block) or with the number of impressions won. (ii) This policy is obtained by solving problem $\hat{\mathcal{P}}(\vec{\beta})$, which is a problem of minimizing a (non-linear) convex function over a polyhedral feasible region, and is therefore efficiently solvable.

2.1.2 A Near-Optimal Solution to Problem $\mathcal{P}_{static}(\vec{M}, \alpha)$

We know that the policy $\pi(\vec{\beta})$ is optimal for Problem $\mathcal{P}^E_{static}(\vec{\beta})$. Given this, we now proceed to find a suitable vector, $\vec{\beta}$, such that $\pi(\vec{\beta})$ is feasible for our original problem $\mathcal{P}_{static}(\vec{M}, \alpha)$ and has a near-optimal expected cost for that problem. The basic intuition in identifying the "correct" $\vec{\beta}$ to use is the following: Recall that, in Problem $\mathcal{P}^E_{static}(\vec{\beta})$, $\vec{\beta}$ is the vector of the expected number of impressions won (for the different campaigns). Notice that, under any feasible policy for $\mathcal{P}_{static}(\vec{M}, \alpha)$, the vector of the expected number of impressions won is larger (componentwise) than $\alpha\vec{M}$. Thus, any feasible policy for $\mathcal{P}_{static}(\vec{M}, \alpha)$ is also feasible for $\mathcal{P}^E_{static}(\alpha\vec{M})$; the converse is, of course, not true. The idea then is to find the optimal policy for $\mathcal{P}^E_{static}(\vec{\beta})$ for a suitably picked "padded vector"

$\vec{\beta} \geq \alpha \vec{M}$. That is, we need to add a "safety cushion" to $\alpha \vec{M}$ and target, in expectation, the resulting inflated vector. Since, for any expected target vector $\vec{\beta}$, the optimal policy $\pi(\vec{\beta})$ is a state independent policy, the number of impressions won for each campaign c over every time block (s, i) has a Binomial distribution. Using the optimal policy for $\mathcal{P}^E_{static}(\vec{\beta})$ from Theorem 1, we can see that the mean of this distribution is $\Delta\, \hat{p}^{\vec{\beta}}_{s,i,c}$, where $\hat{p}^{\vec{\beta}}_{s,i,c} := \sum_{l \in \mathcal{L}_c} q_{i,l} \hat{y}^*_{s,i,l,c}(\vec{\beta})$ and that its variance is $\Delta\, \hat{p}^{\vec{\beta}}_{s,i,c}(1 - \hat{p}^{\vec{\beta}}_{s,i,c})$. Given the large number of time slots over which impressions arrive, this Binomial distribution is approximately normally distributed. Thus, the total number of impressions delivered to campaign c is approximately normally distributed with mean $\sum_{s=1}^{K_c} \sum_{i \in \mathcal{I}} \Delta\, \hat{p}^{\vec{\beta}}_{s,i,c}$ and variance $\sum_{s=1}^{K_c} \sum_{i \in \mathcal{I}} \Delta\, \hat{p}^{\vec{\beta}}_{s,i,c}(1 - \hat{p}^{\vec{\beta}}_{s,i,c})$. Therefore, the probabilistic guarantee associated with campaign i requires us to ensure that $\sum_{s=1}^{K_c} \sum_{i \in \mathcal{I}} \Delta\, \hat{p}^{\vec{\beta}}_{s,i,c} - z_\alpha \sqrt{\sum_{s=1}^{K_c} \sum_{i \in \mathcal{I}} \Delta\, \hat{p}^{\vec{\beta}}_{s,i,c}(1 - \hat{p}^{\vec{\beta}}_{s,i,c})} = M_c,\ c \in \mathcal{C}$, where $z_\alpha := \Phi_N^{-1}(\alpha)$ and Φ_N denotes the standard normal cumulative distribution function. Since the probability that an impression arrives (in any location) in a given time slot is extremely small, the value of $1 - \hat{p}^{\vec{\beta}}_{s,i,c} \approx 1$ for any $\vec{\beta}$. This implies that we need to find $\vec{\beta}$ such that $\sum_{s=1}^{K_c} \sum_{i \in \mathcal{I}} \Delta\, \hat{p}^{\vec{\beta}}_{s,i,c} - z_\alpha \sqrt{\sum_{s=1}^{K_c} \sum_{i \in \mathcal{I}} \Delta\, \hat{p}^{\vec{\beta}}_{s,i,c}} = M_c$. From the definitions of $\hat{y}^*_{s,i,l,c}(\vec{\beta})$ (which is an optimal, and therefore feasible, solution to $\hat{\mathcal{P}}(\vec{\beta})$) and $\hat{p}^{\vec{\beta}}_{s,i,c}$, we know that $\sum_{s=1}^{K_c} \sum_{i \in \mathcal{I}} \Delta\, \hat{p}^{\vec{\beta}}_{s,i,c} = \sum_{s=1}^{K_c} \sum_{i \in \mathcal{I}} \Delta\ \sum_{l \in \mathcal{L}_c} q_{i,l} \hat{y}^*_{s,i,l,c}(\vec{\beta}) = \beta_c$. Thus, the above equation is the same as $\beta_c - z_\alpha \sqrt{\beta_c} = M_c$, solving which we obtain $\beta_c = \frac{2M_c + z_\alpha^2 + \sqrt{(2M_c + z_\alpha^2)^2 - 4M_c^2}}{2}$. In practice, the number of impressions M_c required by any campaign c is of the order of thousands. Also, the platform's probabilistic guarantee α is typically in the range $[0.95, 0.99]$. In such settings, the above expression approximates to $\beta_c = M_c + z_\alpha \sqrt{M_c}$. This has a nice interpretation: The quantity $z_\alpha \sqrt{M_c}$ is the safety cushion for campaign c, similar to the notion of *safety stock* in the newsvendor model. We now show that the policy $\pi(\vec{\beta})$, where β_c is chosen as above (for every c), is near-optimal to $\mathcal{P}_{static}(\vec{M}, \alpha)$.[6]

For $c \in \mathcal{C}$, let $\beta^*_c = M_c + z_\alpha \sqrt{M_c}$. We will use the policy $\pi(\vec{\beta^*})$ as our solution to $\mathcal{P}_{static}(\vec{M}, \alpha)$. Theorem 2 below provides a bound on the performance of this policy. Let $\text{Opt}(\vec{M}, \alpha)$ denote the optimal cost in Problem $\mathcal{P}_{static}(\vec{M}, \alpha)$. Let $\gamma = \max_{c \in \mathcal{C}} \left[\frac{\beta^*_c}{\alpha M_c}\right] = \max_{c \in \mathcal{C}} \left[\frac{1}{\alpha}\left(1 + \frac{z_\alpha}{\sqrt{M_c}}\right)\right]$. Define

$$\psi_{i,l}(x) = \frac{x f'_{i,l}(x)}{f_{i,l}(x)},\ \forall\, i \in \mathcal{I},\ l \in \mathcal{L},\ x \in [0, 1].$$

Since $\alpha \in [0.95, 0.99]$ and M_c is of the order of 100,000, we have $\gamma x_{max} \ll 1$. We proceed under the assumption that $\gamma x_{max} < 1$, and define $\overline{\psi} = \max\limits_{x \in [0, \gamma x_{max}]} \max\limits_{i \in \mathcal{I}, l \in \mathcal{L}} \psi_{i,l}(x)$. Assume that $\overline{\psi} < \frac{\gamma}{\gamma - 1}$: This assumption is easily satisfied in realistic problem instances. For example, in the discussion following Theorem 2, we will see that $\overline{\psi}$ is an order of magnitude smaller than $\frac{\gamma}{\gamma - 1}$.

Theorem 2. *The ratio of the expected cost of policy $\pi(\vec{\beta^*})$ to the optimal cost in Problem $\mathcal{P}_{static}(\vec{M}, \alpha)$ is bounded from above as follows:*

$$\frac{f^{\pi(\vec{\beta^*})}}{Opt(\vec{M}, \alpha)} \le \frac{1}{1 - \frac{(\gamma - 1)}{\gamma}\overline{\psi}}. \tag{2.9}$$

A quantity of immediate interest is the magnitude of the bound $\frac{1}{1 - \frac{(\gamma-1)}{\gamma}\overline{\psi}}$ on realistic problem instances. We shed light on this through the following realistic example. Let $M_c = 150,000$ for all c, $x_{max} = 0.05$ and $\alpha = 0.99$. Then, $\gamma = 1.016$. For the win-curves estimated in Chapter 3 based on real data, we get $\overline{\psi} = 1.42$. Substituting $\gamma = 1.016$ and $\overline{\psi} = 1.42$ in the right-hand side of (2.9), the bound $\frac{1}{1 - \frac{(\gamma-1)}{\gamma}\overline{\psi}}$ equals 1.023. That is, the total cost corresponding to the state-independent policy $\pi(\vec{\beta^*})$ is at most 2.3% higher than that of an optimal policy for $\mathcal{P}_{static}(\vec{M}, \alpha)$.

[6] It is easy to see that our analysis can naturally be extended to one in which the platform provides a campaign-specific probabilistic guarantee α_c. In this case, the value of the "padded demand" β_c will be computed using $\beta_c = M_c + z_{\alpha_c}\sqrt{M_c}$.

We now examine the intuition behind the two parameters γ and $\overline{\psi}$ in (2.9):

- As far as the performance bound of our policy is concerned, the quantity of immediate interest is the extent to which we over procure. This quantity is represented by γ: For a campaign $c \in \mathcal{C}$, the ratio $\frac{M_c + z_\alpha \sqrt{M_c}}{\alpha M_c}$ captures the (percentage) additional number of impressions needed by our policy (to ensure a probabilistic guarantee of α) over the amount αM_c needed by the relaxed problem $\mathcal{P}^E(\alpha \vec{M})$ (which provides a lower bound on the optimal cost) for campaign c. Thus, $\gamma = \max_{c \in \mathcal{C}} \left[\frac{M_c + z_\alpha \sqrt{M_c}}{\alpha M_c} \right]$ is the maximum additional percentage of impressions needed by our policy. For example, if $\gamma = 1.05$, then our policy procures 5% additional impressions.

- Note that $\psi_{i,l}(x) = \frac{f'_{i,l}(x)}{f_{i,l}(x)/x}$. The numerator of the expression on the right-hand-side of the above equation is the marginal expected cost at win-probability x while the denominator is the marginal expected cost at win-probability x if the expected cost function were a linear function with a slope $f_{i,l}(x)/x$. Thus, for a win-probability of x, the ratio $\psi_{i,l}(x)$ compares the (true) marginal cost relative to the (hypothetical) marginal cost if the expected cost function were linear. In other words, $\psi_{i,l}(x)$ is a measure of how fast the expected cost function changes relative to a (hypothetical) linear cost function. Thus, $\overline{\psi} = \max_{x \in [0, \gamma x_{max}]} \max_{i \in \mathcal{I}, l \in \mathcal{L}} \psi_{i,l}(x)$
  is the maximum rate at which the expected cost function increases relative to a linear cost function. For example, if $\overline{\psi} = 1.4$, then the maximum rate at which the expected cost function increases, relative to a linear expected cost function, is 40%.

The following two results highlight the structure of our policy with respect to its win probabilities, i.e., the values $\hat{x}^*_{s,i,l}(\vec{\beta}^*), s \in \mathcal{S}, i \in \mathcal{I}, l \in \mathcal{L}$.

Proposition 1. *For any location $l \in \mathcal{L}$, the target win-probabilities for location l in our policy $\pi(\vec{\beta}^*)$ decrease with time in the following sense: $\hat{x}^*_{s,i,l}(\vec{\beta}^*) \geq \hat{x}^*_{s+1,i,l}(\vec{\beta}^*),\ i \in \mathcal{I},\ 1 \leq s \leq S-1.$*

Proposition 2. *For any location $l \in \mathcal{L}$, any time period $s \in \mathcal{S}$ and any time-block $i \in \mathcal{I}$, the target win-probabilities in our policy $\pi(\vec{\beta}^*)$ decrease with an increase in the arrival probability $q_{i,l}$. That is, as $q_{i,l}$ increases, the win-probability $\hat{x}^*_{s,i,l}(\vec{\beta}^*)$ decreases.*

Remark 1: Since the derivation of our policy $\pi(\vec{\beta}^*)$, which has a closed-form expression, uses the Normal approximation to the c.d.f. of a Binomial distribution, it is possible that this policy slightly violates the probabilistic guarantee that every campaign c is delivered M_c impressions with probability α. This infeasibility can be easily avoided if closed-form expressions are not required. In this case, we can numerically compute a padded vector $\vec{\tilde{\beta}}$ such that policy $\pi(\vec{\tilde{\beta}})$ satisfies the probabilistic guarantee. Such a policy is also endowed with the cost performance guarantee stated in Theorem 2 with the change that γ be replaced by $\tilde{\gamma} := \max_{c \in \mathcal{C}} \left[\frac{\tilde{\beta}_c}{\alpha M_c} \right]$. ■

Remark 2: We now explain how our analysis thus far continues to hold when we relax the assumption that the arrivals of impressions are independent across time slots.

- The results that (i) problem $\mathcal{P}^E_{static}(\alpha \vec{M})$ is a relaxation of problem $\mathcal{P}_{static}(\vec{M}, \alpha)$ and (ii) problem $\mathcal{P}^E_{static}(\vec{\beta})$ is equivalent to problem $\hat{\mathcal{P}}(\vec{\beta})$, continue to hold for an arbitrary stochastic process of impression arrival (after making obvious modifications in the notation).

- The derivation in Theorem 1 of an optimal policy for problem $\mathcal{P}^E_{static}(\vec{\beta})$ by solving $\hat{\mathcal{P}}(\vec{\beta})$ and the fact that this policy is a state-independent policy also continue to hold for an arbitrary stochastic process of impression arrival.

- The next step of our analysis was the computation of the padded vector $\vec{\beta}^*$ such that the optimal policy for problem $\mathcal{P}^E_{static}(\vec{\beta}^*)$ is feasible for $\mathcal{P}_{static}(\vec{M}, \alpha)$. In the absence of the independence assumption, we lose the simple, closed-form expression for $\vec{\beta}^*$. In the literature, Markov Modulated Poisson processes have been used as a popular model to capture dependence in continuous time (see, e.g., Song and Zipkin 1993, Abhyankar and Graves 2001). In our discrete-time setting, this is best modeled by a discrete-time Markov chain which can potentially change states at the beginning of every time block. This state would then dictate the arrival rate of impressions. It is common to assume that this state is observable (for example, the outcome of a certain sports event affecting the arrival rate of impressions of a certain type). In this setting, similar to Remark 1 above, we can numerically compute a padded vector $\vec{\tilde{\beta}}$. Here too, the cost performance guarantee remains as stated in Theorem 2 with the change that γ be replaced by $\tilde{\gamma} := \max_{c\in\mathcal{C}} \left[\frac{\tilde{\beta}_c}{\alpha M_c}\right]$. ■

2.2 Implementation Issues and Policy Recommendations

When an impression arrives, it is auctioned at a mobile ad-exchange in real time. For the platform to participate in the auction, it must typically respond with a bid in less than 100 milliseconds (Downey, 2012). Once all legal bids are in, the ad exchange decides the winner and displays the winner's ad on the impression. The stringent time requirement for the platform's response (i.e., bid) necessitates the adoption of a specific architecture for the procurement and allocation process; real-world platforms typically use a "master-slave" bidding architecture. When advertisers initiate campaigns, their requests are received by a master server, which manages the campaigns by first decomposing the tasks involved in fulfilling the campaigns and then assigning responsibilities to multiple slave servers. The use of multiple slave servers ensures that a (potential) technical failure of one server does not stall the procurement process of all the campaigns. Decomposition ensures that each slave server is largely independent, i.e., it has its own budget and can bid independently, thus saving time by avoiding the need to coordinate with the other slave servers.

One effective decomposition approach is to partition the campaigns across locations: Each slave server is assigned the responsibility of procuring impressions from a set of locations; for each location, the number of impressions to procure, the time limit, and the budget, are provided to the slave server by the master server. The slave servers do not share locations and can, therefore, go about their task without coordinating with the other slave servers. The key question for the master server then is to achieve such a decomposition.

For simplicity of our exposition here, we consider a special case (of the model in Section 2.1) in which all the campaigns start and end at the same time, and the impression-arrival probabilities and the bid curves stay the same through time. The time duration of each campaign is T time slots. We now discuss two decomposition approaches.

2.2.1 An Informed Decomposition Based on Our Near-Optimal Policy

Our policy (Section 2.1.2) lends itself nicely for implementation on this architecture. Recall that this policy requires each location $l \in \mathcal{L}$ to target a constant win-probability $\hat{x}_l^*(\vec{\beta}^*)$ within a time block, to satisfy the platform's probabilistic guarantee α for each campaign. Since we have only one time block here, our policy becomes a *static* policy in this special case. Assume that we have s slave servers and consider a partition of the set of locations $\mathcal{L}$ into s disjoint subsets: $\mathcal{L}^1, \mathcal{L}^2, \ldots, \mathcal{L}^s$. Then, slave server k is assigned the responsibility of winning impressions from the locations in the subset $\mathcal{L}^k, k = 1, 2, \ldots, s$. Although we do not propose any particular partitioning approach, a partition that assigns roughly a similar number of locations to each server or a partition in which each server procures roughly a similar number of impressions should be reasonable in practice. Our policy can now be implemented as follows: For a specific location $l \in \mathcal{L}^k$, slave server k bids $b_l\big(\hat{x}_l^*(\vec{\beta}^*)\big)$ on an impression to target a winning probability of $\hat{x}_l^*(\vec{\beta}^*)$ at that location. The impressions won by the slaves can then be allocated to the various campaigns using the probabilities $\hat{y}_{l,c}^*(\vec{\beta}^*),\ l \in \mathcal{L},\ c \in \mathcal{C}$, specified by the policy. Next, we discuss an intuitive approach for decomposing campaigns across locations.

2.2.2 Greedy Decomposition

Assume that the bid curves $b_l(x)$ can be *strictly ordered* across locations: For any two locations $l, l' \in \mathcal{L}$, exactly one of the following holds: $b_l(x) < b_{l'}(x), \forall x \in (0,1]$, or $b_l(x) > b_{l'}(x),\ \forall x \in (0,1]$. It then follows that each campaign has a unique cheapest location. For campaign $c \in \mathcal{C}$, let $l_c^* \in \mathcal{L}_c$ denote its cheapest location, i.e., $b_{l_c^*}(x) < b_l(x),\ \forall l \in \mathcal{L}_c \backslash \{l_c^*\}, x \in (0,1]$.

In the greedy decomposition approach, each campaign $c \in \mathcal{C}$ procures all its impressions from its cheapest location l_c^*. Since the set of locations for the different campaigns may overlap, it is possible for a given location to be the cheapest location for multiple campaigns. Let $\mathcal{C}_l^*$ be the set of campaigns for which location l is the cheapest location, i.e., $\mathcal{C}_l^* = \{c \in \mathcal{C} : l_c^* = l\}$. Thus, all the impressions needed for the campaigns in $\mathcal{C}_l^*$ will be sourced from location l. Recall from our analysis in Section 2.1.2 that the expected number of impressions needed for a campaign $c \in \mathcal{C}$ to meet the platform's probabilistic guarantee is $\beta_c^* = M_c + z_\alpha \sqrt{M_c}$. Thus, the aggregate number of impressions to be procured from location $l \in \mathcal{L}$ is $M_l = \sum_{c \in \mathcal{C}_l^*} \beta_c^*$. Using the same notation as in the previous subsection, suppose that slave server k is responsible for winning impressions from the locations in the subset $\mathcal{L}^k, k = 1, 2, \ldots, s$. Then, if slave server k uses a static policy, it targets a win-probability of $x_l = \frac{M_l}{T q_l}$ for a specific location $l \in \mathcal{L}^k$.

2.2.3 A Comparison of Decomposition and Procurement Approaches

By using either of the decomposition approaches discussed above, suppose that a slave server is assigned the task of procuring M_l impressions from location $l \in \mathcal{L}$. Let $j_{t,l}$ be the number of impressions procured from location l by the end of time slot $t-1$. Thus, at the beginning of time slot t, the remaining number of impressions to be procured from location l is $M_l - j_{t,l}$ while $q_l(T-t)$ is the expected number of impressions yet to arrive from location l over the campaign time-duration T. Accordingly, the server targets a win-probability of $x_{l,t} = \min\left\{\frac{M_l - j_{t,l}}{q_l(T-t)}, 1\right\}$ at location l; i.e., it bids $b_l(x_{l,t})$ in state $j_{t,l}$. Clearly, the server's bidding intensity in this policy changes with the state. For example, if the remaining number of impressions to be procured approaches the expected number of impressions to arrive, then the bids become aggressive. This policy is based on our observations at Cidewalk.

The preceding discussion identifies two approaches each for the decomposition of campaigns across slaves (namely, *Informed* and *Greedy*) and a server's procurement policy (namely, *Static* and *Reactive*). As a result, we have the following four approaches: (i) *Informed-Static* (IS) – The informed approach (I) for decomposition followed by the static procurement policy (S). (ii) *Informed-Reactive* (IR) – The informed approach (I) for decomposition followed by the reactive procurement policy (R). (iii) *Greedy-Static* (GS) – The greedy approach (G) for decomposition followed by the static procurement policy (S). (iv) *Greedy-Reactive* (GR) – The greedy approach (G) for decomposition followed by the reactive procurement policy (R). We now use numerical experiments to examine these approaches for different practical settings of the platform's problem and then offer our recommendations.

Simulation and Recommendations

We begin by specifying the default setting of the simulation. To assess the impact of a specific parameter, we will vary only that parameter while retaining the remainder of the default setting.

There are 4 campaigns and 4 locations, i.e., $\mathcal{C} = \{1, 2, 3, 4\}$ and $\mathcal{L} = \{1, 2, 3, 4\}$. Each campaign needs 3000 impressions within a time-duration of $T = 10^6$ time slots. Thus, $M_c = 3000, \forall c \in \mathcal{C}$. For each of the four locations, the probability that an impression will arrive in a time slot is 0.05, i.e., $q_l = 0.05,\ \forall l \in \mathcal{L}$. The platform's probabilistic guarantee α is 0.99. We choose a simple polynomial form for a bid curve that makes it easy to vary the heterogeneity in the bid curves across locations: For location $l \in \mathcal{L}$, the bid curve is $b_l(x) = \left(1 + \frac{(a-1)l}{4}\right)x^4$, where $a \geq 1$ is a parameter; thus, as a increases, the locations become increasingly heterogeneous. We use $a = 2$ as a default. Also, notice that the cost of winning impressions at the locations is in the order of their index l, with location 1 being the cheapest. The locations of interest for the four campaigns, i.e., the sets $\mathcal{L}_c, c \in \mathcal{C}$, are specified by a 0-1 "association" matrix $\langle c, l\rangle, c \in \mathcal{C}, l \in \mathcal{L}$: the entry $(c, l) = 1$, if campaign c is interested in location l; 0, otherwise. Thus, for campaign $c \in \mathcal{C}$, the set $\mathcal{L}_c$ consists of the locations with entry 1 in the row of the matrix corresponding to campaign c. Four different association matrices are illustrated in Figure 2.1; we will soon discuss the interpretations of these matrices. We use the association matrix A_2 as our default.

$$A_1 = \begin{array}{c} \\ c_1 \\ c_2 \\ c_3 \\ c_4 \end{array} \begin{array}{c} \begin{array}{cccc} l_1 & l_2 & l_3 & l_4 \end{array} \\ \begin{bmatrix} 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix} \end{array} \quad A_2 = \begin{array}{c} \\ c_1 \\ c_2 \\ c_3 \\ c_4 \end{array} \begin{array}{c} \begin{array}{cccc} l_1 & l_2 & l_3 & l_4 \end{array} \\ \begin{bmatrix} 0 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix} \end{array} \quad A_3 = \begin{array}{c} \\ c_1 \\ c_2 \\ c_3 \\ c_4 \end{array} \begin{array}{c} \begin{array}{cccc} l_1 & l_2 & l_3 & l_4 \end{array} \\ \begin{bmatrix} 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix} \end{array} \quad A_4 = \begin{array}{c} \\ c_1 \\ c_2 \\ c_3 \\ c_4 \end{array} \begin{array}{c} \begin{array}{cccc} l_1 & l_2 & l_3 & l_4 \end{array} \\ \begin{bmatrix} 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix} \end{array}$$

Low Overlap | Medium Overlap | Medium Overlap | High Overlap

Figure 2.1. Matrices representing increasing levels of overlap in the cheapest locations for the campaigns.

We now discuss three important properties that we wish to investigate to compare the four approaches defined above, namely IS, IR, GS, and GR.

- **Overlap in the Locations Across Campaigns:** The extent of overlap in the locations across the campaigns is difficult to characterize succinctly: For two campaigns, not only is the *number* of overlapping locations important but also are the *identities* of these locations. Also, both the number and the identities of the overlapping locations are likely to vary across pairs of campaigns. We use four campaign-location matrices to vary the extent of overlap across campaigns; Figure 2.1 defines these four matrices. Recall that the locations are indexed in the increasing order of their bid curves $b_l(\cdot)$. For the first matrix A_1, the cheapest location for campaign c_i is $l_i, i = 1, 2, 3, 4$; thus, there is no overlap in the cheapest locations for the four campaigns. On the other hand, for the matrix A_4, the cheapest location is the same (namely, l_3) for campaigns c_1, c_2, and c_3; thus, matrix A_4 corresponds to a setting with a high overlap. The remaining two matrices (A_2 and A_3) are intermediate settings that represent a medium amount of overlap.

- **Total Demand Over the Campaigns:** Broadly, the higher the total demand, the higher should be the platform's bidding intensity. We consider the following set of values for the demand of each campaign: $\{1000, 2000, 3000, \ldots, 7000\}$.

- **Heterogeneity in the Cost of Procuring Impressions Across Locations:** An increase of the heterogeneity in the bid curves across locations leads to the platform increasingly preferring certain locations over others. We use the following set of values of the parameter a in the expression for the bid curves (defined above) to vary the heterogeneity in the cost of procuring impressions across locations: $\{1, 2, 3, \ldots, 10\}$.

We now discuss our main observations from the simulation.

Effect of Overlap in the Locations: Figure 2.2(a) compares the four approaches as the extent of overlap in the campaign locations increases. Clearly, when there is little or no overlap, then the decomposition method does not have a significant bearing on the total cost: both methods prefer procuring from the cheapest locations of the respective campaigns. Moreover, the cheapest locations of the campaigns are largely distinct, and it is therefore possible to source the required impressions from these locations. Thus, the static approaches (IS and GS) incur nearly the same costs. Similarly, the reactive approaches (IR and GR) too incur similar, but relatively higher, costs. As the overlap increases, the benefit of informed decomposition becomes apparent and the static approaches become significantly superior. With an increase in the overlap, the set of cheap locations available for procuring the impressions shrinks. Consequently, the number of impressions required from these cheap locations progressively increases, which in turn increases the target win-probabilities at these locations and makes the reactive policy compete better with the static policy (for instance, when the target win-probabilities are close to 1, it is easy to see that the costs of the static and reactive policies are similar for a given decomposition method). The same effect vis-à-vis the comparison of the static and reactive policies can be seen in the two greedy approaches as well.

Under a very high overlap, it may become expensive to procure impressions. In such a situation, it is possible that the reactive policy (IR) might perform marginally better than the static policy (IS), for the following reason: In the latter, we procure $M_c + z_\alpha\sqrt{M_c}$ impressions for each campaign c (see Section 2.1.2), while in the former we procure M_c impressions. When impressions are expensive to procure, then the advantage that the static policy has – in terms of being derived by solving a global optimization problem, namely $\hat{\mathcal{P}}(\vec{\beta})$ – can sometimes get negated by the additional cost of procuring the extra $z_\alpha\sqrt{M_c}$ impressions.

Effect of Increase in Demand: Figure 2.2(b) shows the impact of increase in the demand of impressions faced by the campaigns. When demand is low, it is possible to procure most of the impressions needed for a campaign from its cheapest location. Consequently, both the informed and greedy decompositions behave in a similar manner and the decomposition method does not have a significant effect on the total cost. Hence, the static approaches (IS and GS) cost nearly the same and so too do the reactive approaches (IR and GR), with the static approaches being relatively better. As the demand of impressions for a campaign increases, both the use of multiple locations to satisfy the demand as well as a careful division of this demand across these locations become important, which is

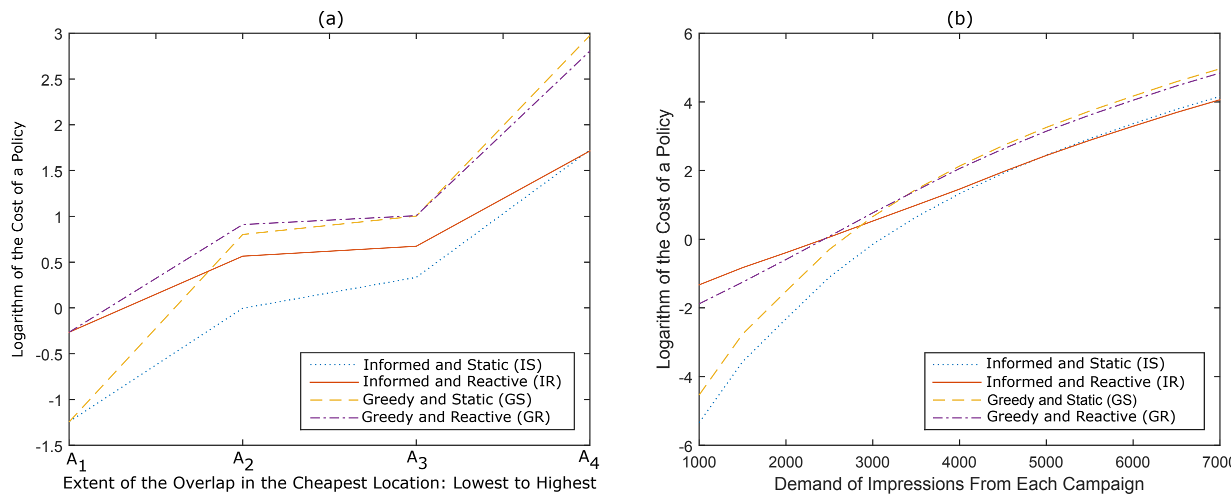

Figure 2.2. Effect of overlap in the cheapest locations across campaigns and increase in the demand of the campaigns.

exactly what informed decomposition achieves. Consequently, the benefit from informed decomposition progressively increases with the demand for impressions.

The target win-probabilities of the campaigns (at their respective sets of preferred locations) increase with demand and result in an improvement in the relative performance of the reactive policy. Recall from Section 2.1.2 that the target win-probability of our policy is computed after inflating the required number of impressions (by the additional amount $z_\alpha\sqrt{M_c}$) to ensure that the platform's probabilistic guarantee is met. The reactive policy, on the other hand, does not need this additional amount. Therefore, when the demand for impressions is high (and hence they are expensive to procure), the reactive policy may perform marginally better than the static policy.

Effect of Bid-Curve Heterogeneity: An increase in the heterogeneity in the bid curves (i.e., the cost of procuring impressions) across geographical locations naturally leads to the campaigns increasingly favoring their cheapest locations. Therefore, *the costs corresponding to the greedy and informed decompositions come closer as bid-curve heterogeneity increases.* To isolate this effect, the left-hand-side subfigure of Figure 2.3 plots two ratios: (i)

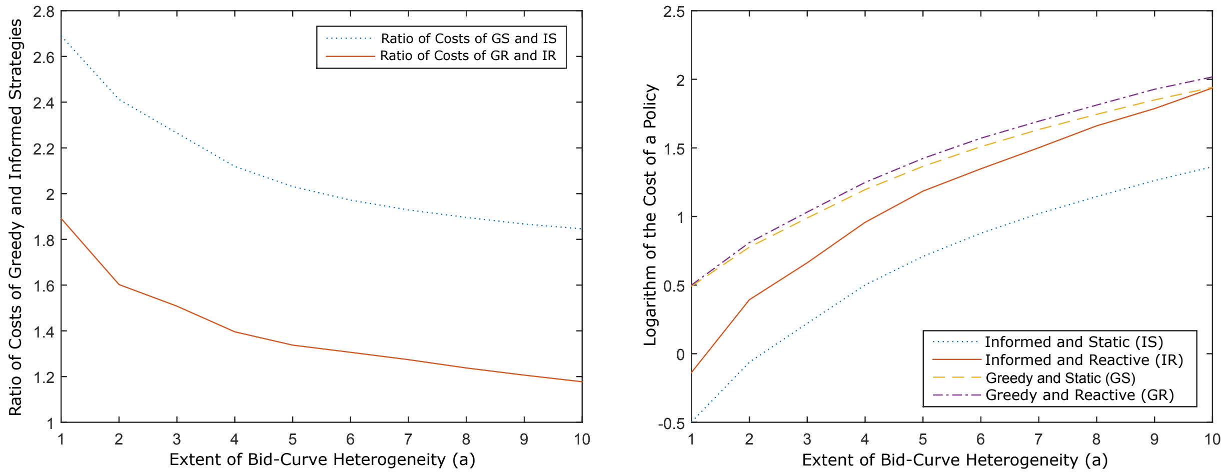

Figure 2.3. Effect of bid-curve heterogeneity.

The ratio $\frac{\text{Cost}_{\text{GS}}}{\text{Cost}_{\text{IS}}}$ of the costs corresponding to the greedy and informed decompositions when the procurement policy is static and (ii) The ratio $\frac{\text{Cost}_{\text{GR}}}{\text{Cost}_{\text{IR}}}$ of the costs corresponding to greedy and informed decomposition when the procurement policy is reactive. Both these ratios decrease as bid-curve heterogeneity increases.

The right-hand-side subfigure of Figure 2.3 shows the costs of the four approaches. As expected, the combination of greedy decomposition and reactive procurement (i.e., the GR approach) is the worst. To improve from the GR approach, there are three possible choices: (i) Shift to informed decomposition, while continuing to use the reactive procurement (approach IR). (ii) Shift to static procurement, while continuing to use greedy decomposition (approach GS). (iii) Change both the decomposition method and procurement policy (approach IS). The figure shows an interesting *compounding effect*: The reduction in cost realized by improving both the the decomposition method and procurement policy (i.e., moving from GR to IS) is greater than the sum of the reductions achieved by improving either the decomposition alone or the procurement policy alone.

CHAPTER 3

ANALYSIS OF A DYNAMIC MODEL OF CAMPAIGN ARRIVALS[1]

In this chapter, we generalize our analysis of the first chapter and consider a dynamic model of campaign arrivals. First, we describe the model. Recall that in Section 2.1, we assumed that one campaign of each type arrives in period 1 (the first period in a finite planning horizon) and that there are no further arrivals of campaigns. In contrast, we now consider an infinite horizon[2] and assume that, at the beginning of every period $s \in \mathbb{N}$, one campaign of each type arrives (a non-arrival is modeled as a campaign with a requirement of zero impressions). For expositional convenience, we assume that all campaigns have the same time duration of K time periods. Let (s, c) denote the type-c campaign that arrives in period s and let $\tilde{M}_{s,c}$ be the random variable representing the number of impressions required by campaign (s, c). We allow any correlation structure between these random variables, for a given period, s. The joint distribution of these random variables is known. Observe that we allow the possibility that campaigns of all types might not arrive in every period; for example, if $\mathbb{P}(\tilde{M}_{s,c} = 0) > 0$, there is a positive probability that a type c campaign does not arrive in period s. (Note: Our assumption that at most one campaign of any given type can arrive in a period is made exclusively for notational simplicity.) The random vector $\tilde{\mathbf{M}}_s$ of requirements corresponding to each of the campaign types in period s is assumed to be *independently and identically distributed across time periods*, that is, the sequence $\{\tilde{\mathbf{M}}_s\}$ is i.i.d.. Let $M_c = \mathbb{E}[\tilde{M}_{s,c}]$ be the expected requirement of a type-c campaign, $c \in \mathcal{C}$.

---

[1]Republished with permission of INFORMS, from "Procurement Policies for Mobile-Promotion Platforms," M. Aseri, M. Dawande, G. Janakiraman, and V. Mookerjee, forthcoming in Management Science, October 2017; permission conveyed through Copyright Clearance Center, Inc.

[2]Therefore, the set of time periods in the horizon is $\mathcal{S} = \mathbb{N}$ and the set of time slots in the horizon is also $\mathcal{T} = \mathbb{N}$.

Consider any period s and any $i \in \mathcal{I}$. For every $t \in \mathcal{T}_{s,i}$, we define $s(t) = s$; that is, $s(t)$ denotes the period to which slot t belongs. For any $s \leq s(t)$, let $j_{t,(s,c)}$ denote the number of impressions won by campaign (s,c) before the start of time-slot t. Thus, $\vec{j}_t = (j_{t,(s,c)} : s(t) - K + 1 \leq s \leq s(t),\ c \in \mathcal{C})$ denotes the state vector at the beginning of time-slot t; this vector captures the information on the number of impressions won by the various campaigns that have arrived before time slot t but have not ended by that time. Let $\mathcal{J}_t$ denote the set of all feasible state vectors in time-slot t.

A policy π is defined by a set of target win-probability functions $\{x^{\pi}_{t,l} : \mathcal{J}_t \to [0,1],\ t \in \mathcal{T},\ l \in \mathcal{L}\}$ and a set of allocation functions $\{y^{\pi}_{t,l,(s,c)} : \mathcal{J}_t \to [0,1],\ t \in \mathcal{T},\ l \in \mathcal{L},\ s \leq s(t),\ c \in \mathcal{C}\}$. If an impression arises from location l in time-slot t, then $x^{\pi}_{t,l}(\vec{j}_t)$ denotes the win-probability targeted by π for this impression and $y^{\pi}_{t,l,(s,c)}(\vec{j}_t)$ denotes the probability with which policy π wins this impression and allocates it to campaign (s,c).[3] We also assume an upper bound $x_{max} \in (0,1]$ on the target win-probabilities $\{x^{\pi}_{t,l}\}$.

The random outcomes associated with the arrival, the winning, and the allocation of impressions are defined in exactly the same manner as in Section 2.1 using the random variables $\{U^{imp}_t\}$, $\{U^{win}_t\}$ and $\{U^{alloc}_t\}$. Recall also that $\tilde{\mathbf{M}}_s$ denotes the vector of requirements for the various campaign types arriving in period s. For any t, let

$$\vec{U}_t = \left(U^{imp}_1, U^{win}_1, U^{alloc}_1, U^{imp}_2, U^{win}_2, U^{alloc}_2, \ldots, U^{imp}_t, U^{win}_t, U^{alloc}_t, \tilde{\mathbf{M}}_1, \tilde{\mathbf{M}}_2, \ldots, \tilde{\mathbf{M}}_{s(t)}\right)$$

be the history of these random variables until time-slot t (including t) and let $\mathcal{U}_t$ be the set of all possible vectors $\vec{U}_t$. From the definition of $x^{\pi}_{t,l}(\vec{j}_t)$ and $y^{\pi}_{t,l,(s,c)}(\vec{j}_t)$, we note that $x^{\pi}_{t,l}(\vec{j}_t) = \sum_{(s,c): s \leq s(t), c \in \mathcal{C}} y^{\pi}_{t,l,(s,c)}(\vec{j}_t),\ \forall\ \vec{j}_t \in \mathcal{J}_t$. Furthermore, since a type-c campaign is only interested in the set of locations $\mathcal{L}_c$, we have $y^{\pi}_{t,l,(s,c)}(\vec{j}_t) = 0$, if $l \notin \mathcal{L}_c$, $\forall c \in \mathcal{C}$, $\vec{j}_t \in \mathcal{J}_t$.

[3] We note that the decisions $x^{\pi}_{t,l}$ and $y^{\pi}_{t,l,c}$ depend on the history until time slot t only through $\vec{j}_t$. This is because of the assumed independence of the arrival of impressions (campaigns) across time slots (periods).

For any policy π, let

$$f^{\pi} := \lim_{S\to\infty} \frac{\sum_{s=1}^{S} \sum_{i\in\mathcal{I}} \sum_{t\in\mathcal{T}_{s,i}} \sum_{l\in\mathcal{L}} \mathbb{E}_{\vec{U}_{t-1}}\left[q_{i,l} f_{i,l}\left(x^{\pi}_{t,l}\left(\vec{j}^{\pi}_{t}(\vec{U}_{t-1})\right)\right)\right]}{S}$$

denote the expected per-period cost of policy π. Below, we present a mathematical formulation of the platform's problem $\mathcal{P}_{Dyn}(\alpha)$. Here, we will use the notation $\leq_l$ $(<_l)$ to denote smaller (strictly smaller) in the lexicographic order.[4]

$$\mathcal{P}_{Dyn}(\alpha): \begin{cases} \min_{\pi}\ f^{\pi} \\ \text{subject to:} \\ x^{\pi}_{t,l}(\vec{j}) \ = \ \sum_{(s,c):s\leq s(t),c\in\mathcal{C}} y^{\pi}_{t,l,(s,c)}(\vec{j}),\ \forall\, t\in\mathcal{T},\ \vec{j}\in\mathcal{J}_t,\ l\in\mathcal{L}, & (3.1) \\ j^{\pi}_{t+1,(s,c)}(\vec{u}_t) \ = \ j^{\pi}_{t,(s,c)}(\vec{u}_{t-1}) \ + \ \sum_{l\in\mathcal{L}_c} \mathbb{1}(U^{imp}_t = l) \\ \times \mathbb{1}\left(U^{win}_t \leq x^{\pi}_{t,l}\left(\vec{j}^{\pi}_t(\vec{u}_{t-1})\right)\right) \\ \times \mathbb{1}\left(\frac{\sum_{(\hat{s},\hat{c})\ <_l\ (s,c)} y^{\pi}_{t,l,(\hat{s},\hat{c})}\left(\vec{j}^{\pi}_t(\vec{u}_{t-1})\right)}{x^{\pi}_{t,l}\left(\vec{j}^{\pi}_t(\vec{u}_{t-1})\right)} < U^{alloc}_t \leq \frac{\sum_{(\hat{s},\hat{c})\ \leq_l\ (s,c)} y^{\pi}_{t,l,(\hat{s},\hat{c})}\left(\vec{j}^{\pi}_t(\vec{u}_{t-1})\right)}{x^{\pi}_{t,l}\left(\vec{j}^{\pi}_t(\vec{u}_{t-1})\right)}\right) \\ \qquad\qquad \forall\, t\in\mathcal{T},\ \forall\, s\leq s(t),\ c\in\mathcal{C},\ \vec{u}_t\in\mathcal{U}_t, & (3.2) \\ x^{\pi}_{t,l}(\vec{j}),\ y^{\pi}_{t,l,(s,c)}(\vec{j}) \in [0, x_{max}], \\ \qquad\qquad \forall\, t\in\mathcal{T}, s\leq s(t),\ \vec{j}\in\mathcal{J}_t,\ l\in\mathcal{L},\ c\in\mathcal{C}, & (3.3) \\ \mathbb{P}\left[j^{\pi}_{(s+K-1)I\Delta,(s,c)}(\vec{U}_{(s+K-1)I\Delta}) \ \geq \ \tilde{M}_{(s,c)} \mid \vec{U}_{(s-1)I\Delta+1} = \vec{u}\right] \ \geq \ \alpha, \\ \qquad\qquad \forall\, c\in\mathcal{C},\ \forall s\in\mathcal{S},\ \forall\vec{u}\in\mathcal{U}_{(s-1)I\Delta+1}. & (3.4) \end{cases}$$

Equation (3.2) defines the transition of $j^{\pi}_{t,(s,c)}(\vec{u}_{t-1})$ to $j^{\pi}_{t+1,(s,c)}(\vec{u}_t)$. To understand this equation, note that the first indicator variable on its right side takes the value 1 if an impression arrived at location l in time slot t, the second indicator takes the value 1 if this impression is won by policy π, and the third indicator takes the value 1 if the impression is allocated to campaign (s,c). Inequality (3.4) states that the probability that $\tilde{M}_{(s,c)}$ impressions will be delivered to campaign (s,c) over that campaign's duration (which ends in the slot

[4] For any $\hat{s}, s\in\mathcal{S}$ and $\hat{c}, c\in\mathcal{C}$, we say that $(\hat{s},\hat{c}) \leq_l (s,c)$ if either (a) $\hat{s} < s$ or (b) $\hat{s} = s$ and $\hat{c}\leq c$. We say $(\hat{s},\hat{c}) <_l (s,c)$ if $(\hat{s},\hat{c}) \leq_l (s,c)$ and $(\hat{s},\hat{c}) \neq (s,c)$.

$(s+K-1)I\Delta\big)$, conditional on the history of the system at the time when that campaign arrives (i.e., beginning of period s which corresponds to time slot $(s-1)I\Delta+1$), exceeds α.

3.1 Lower Bound on the Optimal Per-Period Cost

Let $\mathcal{P}_{E,Dyn}(\alpha)$ represent the problem in which campaigns arrive dynamically, and the constraint is the following *expectation constraint* as opposed to the probabilistic guarantee in Problem $\mathcal{P}_{Dyn}(\alpha)$: For every (s,c), the expected number of impressions delivered is at least $\alpha\tilde{M}_{s,c}$; this constraint is mathematically represented by inequality (3.8) below. It is easy to see that $\mathcal{P}_{Dyn}(\alpha)$ is a relaxation of $\mathcal{P}_{E,Dyn}(\alpha)$. Thus, we have

$$C^*_{Dyn}(\alpha) \geq C^*_{E,Dyn}(\alpha)\ , \tag{3.5}$$

where $C^*_{Dyn}(\alpha)$ and $C^*_{E,Dyn}(\alpha)$ denote the optimal costs in $\mathcal{P}_{Dyn}(\alpha)$ and $\mathcal{P}_{E,Dyn}(\alpha)$, respectively.

Consider any policy $\pi = (\mathbf{x}^\pi,\ \mathbf{y}^\pi)$ which is optimal (therefore, feasible) for $\mathcal{P}_{E,Dyn}(\alpha)$. Thus,

$$C^*_{E,Dyn}(\alpha) = f^\pi = \lim_{S\to\infty} \frac{\sum_{s=1}^{S}\sum_{i\in\mathcal{I}}\sum_{t\in\mathcal{T}_{s,i}}\sum_{l\in\mathcal{L}}\mathbb{E}_{\vec{U}_{t-1}}\Big[q_{i,l}f_{i,l}\Big(x^\pi_{t,l}\big(\vec{j}^\pi_t(\vec{U}_{t-1})\big)\Big)\Big]}{S}. \tag{3.6}$$

From the feasibility of policy π for problem $\mathcal{P}_{E,Dyn}(\alpha)$, we know that π satisfies the following two relations:

$$\begin{aligned} x^\pi_{t,l}(\vec{j}) &\geq \sum_{s=s(t)-K+1}^{s(t)} \sum_{c\in\mathcal{C}} y^\pi_{t,l,(s,c)}(\vec{j}),\ \forall\, t\in\mathcal{T},\ l\in\mathcal{L},\ \vec{j}\in\mathcal{J}_t, && (3.7)\\ \alpha\tilde{M}_{s,c} &\leq \sum_{\hat{s}=s}^{s+K-1}\sum_{i=1}^{I}\sum_{t\in\mathcal{T}_{\hat{s},i}}\sum_{l\in\mathcal{L}_c} q_{i,l}\mathbb{E}[y^\pi_{t,l,(s,c)}\big(\vec{j}^\pi_t(\vec{U}_{t-1})\big) \mid \vec{U}_{(s-1)I\Delta+1}=\vec{u}] \\ &\quad \forall\, c\in\mathcal{C},\ \forall s\in\mathcal{S},\ \forall \vec{u}\in\mathcal{U}_{(s-1)I\Delta+1}. && (3.8)\end{aligned}$$

Inequality (3.7) holds because, for any time slot t and any location l, the corresponding impression (if one arrives) can be allocated only to campaigns which arrive in the set of time periods starting in $s(t) - K + 1$. Inequality (3.8) holds because, for any campaign (s, c), the expected guaranteed number of impressions (i.e., $\alpha\tilde{M}_{s,c}$) should be generated from impressions arising from any time slot in periods $s, s+1, \dots, s+K-1$ and from any location $l \in \mathcal{L}_c$.

We present two useful definitions next. Let $\tilde{x}^{\pi}_{i,l} = \lim_{S\to\infty} \frac{\sum_{s=1}^{S}\sum_{t\in\mathcal{T}_{s,i}} \mathbb{E}[x^{\pi}_{t,l}(\vec{j}^{\pi}_t(\vec{U}_{t-1}))]}{S\Delta}$ denote the average win-probability targeted from location l over all time slots belonging to block i. Similarly, let

$$\tilde{y}^{\pi}_{i,l,c} \quad = \quad \lim_{S\to\infty} \frac{\sum_{s=1}^{S}\sum_{\hat{s}=s}^{s+K-1}\sum_{t\in\mathcal{T}_{\hat{s},i}} \mathbb{E}[y^{\pi}_{t,l,(s,c)}\big(\vec{j}^{\pi}_t(\vec{U}_{t-1})\big)]}{S\Delta}$$

denote the average probability (averaged over all impressions at location l and over all time slots belonging to block i) of winning an impression for a type-c campaign.

Then, applying Jensen's inequality on $f_{i,l}(\cdot)$ in (3.6), we obtain

$$C^{*}_{E,Dyn}(\alpha) \geq \sum_{i\in\mathcal{I}}\sum_{l\in\mathcal{L}} \left[q_{i,l} f_{i,l}\left(\tilde{x}^{\pi}_{i,l}\right)\right]\Delta.$$

Moreover, it is straight forward to show that inequality (3.7) implies that, for any i and l, we have $\tilde{x}^{\pi}_{i,l} \geq \sum_{c\in\mathcal{C}} \tilde{y}^{\pi}_{i,l,c}$. Similarly, using the definition $M_c = \mathbb{E}[\tilde{M}_{s,c}]$, it is easy to see that inequality (3.8) implies that, for any c, $\alpha M_c \leq \sum_{i\in\mathcal{I}}\sum_{l\in\mathcal{L}_c} q_{i,l}\tilde{y}^{\pi}_{i,l,c}\Delta$.

The three inequalities above, along with (3.5), imply that $C^{*}_{Dyn}(\alpha) \geq G_1(\alpha\mathbf{M})$, where $\mathbf{M} := (M_c : c \in \mathcal{C})$ and, for any $\mathbf{m} \in \mathbb{R}^{C,+}$,

$$\begin{aligned}
G_1(\mathbf{m}) := \quad \min_{\tilde{\mathbf{x}},\tilde{\mathbf{y}}} \quad & \Delta\sum_{i\in\mathcal{I}}\sum_{l\in\mathcal{L}}\Big[q_{i,l} f_{i,l}(\tilde{x}_{i,l})\Big] \\
\text{s.t.} \quad & \tilde{x}_{i,l} \geq \sum_{c\in\mathcal{C}} \tilde{y}_{i,l,c} \;\forall\, i\in\mathcal{I}, l\in\mathcal{L}, \\
& m_c = \sum_{i\in\mathcal{I}}\sum_{l\in\mathcal{L}_c} q_{i,l}\tilde{y}_{i,l,c}\Delta \;\forall c\in\mathcal{C}, \\
& \tilde{x}_{i,l} \in [0, x_{max}] \;\forall\, i\in\mathcal{I}, l\in\mathcal{L}, \text{ and} \\
& \tilde{y}_{i,l,c} \in [0, x_{max}] \;\forall\, i\in\mathcal{I}, l\in\mathcal{L}, c\in\mathcal{C}.
\end{aligned}$$

Let the math program above be referred to as $\mathcal{P}_1(\mathbf{m})$: this program represents a single-period problem with a vector of expected requirements of $\mathbf{m}$ for C campaigns, one of each type in $\mathcal{C}$.

3.2 Rolling-Horizon Policy with a Horizon Length of One Period

In this section, we define a rolling-horizon policy with a horizon length of one period. Let us denote this policy by π_1. For every period s, this policy computes its decisions for that period at the beginning of that period by solving a single-period optimization problem defined below. We will show that this policy is a feasible solution to Problem $\mathcal{P}_{Dyn}(\alpha)$.

Consider any $s \in \mathbb{N}$. For any $\hat{s} \in \{s-K+1, s-K+2, \ldots, s\}$ and any $c \in \mathcal{C}$, let $\tilde{m}_{\hat{s},c}$ denote the realization of $\tilde{M}_{\hat{s},c}$ and let $\hat{m}_{\hat{s},c} = \frac{\tilde{m}_{\hat{s},c}}{K} + z_\alpha \sqrt{\frac{\tilde{m}_{\hat{s},c}}{K}}$. This quantity can be interpreted as the per-period requirement of campaign $(\hat{s}, c)$ plus a carefully chosen safety cushion. Also, for every $s \in \mathbb{N}$ and $c \in \mathcal{C}$, define $\bar{m}_{s,c} = \sum_{\hat{s}=s-K+1}^{s} \hat{m}_{\hat{s},c}$, where $\hat{m}_{\hat{s},c} := 0$ if $\hat{s} \leq 0$. Let $\tilde{\mathbf{x}}^*$ and $\tilde{\mathbf{y}}^*$ denote the optimal solution to the single-period problem $\mathcal{P}_1(\bar{\mathbf{m}}_\mathbf{s})$, where $\bar{\mathbf{m}}_\mathbf{s} = (\bar{m}_{s,c} : c \in \mathcal{C})$. Then, the decisions of policy π_1 in period s are defined as follows:

$$x_{t,l}^{\pi_1} = \tilde{x}_{i,l}^* \quad \forall\, i \in \mathcal{I}, t \in \mathcal{T}_{s,i}, l \in \mathcal{L}, \text{ and} \tag{3.9}$$

$$y_{t,l,(\hat{s},c)}^{\pi_1} = \tilde{y}_{i,l,c}^* \left(\frac{\hat{m}_{\hat{s},c}}{\bar{m}_{s,c}} \right) \forall\, c \in \mathcal{C}, i \in \mathcal{I}, t \in \mathcal{T}_{s,i}, l \in \mathcal{L}, \hat{s} \in \{s-K+1, s-K+2, \ldots, s\} \tag{3.10}$$

Equation (3.10) ensures that the procured impressions of type c are allocated to the campaigns of that type in the proportion of their per-period requirements $\hat{m}_{\hat{s},c}$.

Feasibility of Policy π_1 to Problem $\mathcal{P}_{Dyn}(\alpha)$: For any $\hat{s} \in \{s-K+1, s-K+2, \ldots, s\}$, the number of impressions won for campaign $(\hat{s}, c)$ in block i of period s is binomially distributed with parameters Δ and $\sum_{l \in \mathcal{L}} q_{i,l} y_{t,l,(\hat{s},c)}^{\pi_1} = \sum_{l \in \mathcal{L}} q_{i,l} \tilde{y}_{i,l,c}^* \left(\frac{\hat{m}_{\hat{s},c}}{\bar{m}_{s,c}} \right)$. Since Δ is large and $q_{i,l}$ is very small, this binomial distribution can be approximated by a Normal distribution with mean and variance both equal to $\Delta \sum_{l \in \mathcal{L}} q_{i,l} \tilde{y}_{i,l,c}^* \left(\frac{\hat{m}_{\hat{s},c}}{\bar{m}_{s,c}} \right)$. Thus, the number of impressions won for campaign $(\hat{s}, c)$ from all the time blocks of period s is approximately normally distributed with mean and variance both equal to $\Delta \sum_{i \in \mathcal{I}} \sum_{l \in \mathcal{L}} q_{i,l} \tilde{y}^*_{i,l,c} \left(\frac{\hat{m}_{\hat{s},c}}{\bar{m}_{s,c}} \right)$. Observe that the feasibility of $\tilde{\mathbf{x}}^*$ and $\tilde{\mathbf{y}}^*$ to the single-period problem $\mathcal{P}_1(\bar{\mathbf{m}}_\mathbf{s})$ implies the following:

$$\Delta \sum_{i \in \mathcal{I}} \sum_{l \in \mathcal{L}_c} q_{i,l} \tilde{y}^*_{i,l,c} = \bar{m}_{s,c}. \tag{3.11}$$

Therefore, from (3.10) and (3.11), the mean (variance) of the number of impressions won for campaign $(\hat{s}, c)$ from all the time blocks of period s is equal to (approximately equal to) $\hat{m}_{\hat{s},c}$. Thus, the number of impressions won for any campaign $\hat{s} \in \mathcal{S}_{camp}$ during its K-period duration $\{\hat{s}, \hat{s}+1, \ldots, \hat{s}+K-1\}$ is approximately normally distributed with mean and variance $K\hat{m}_{\hat{s},c}$ since the mean and variance of a sum of K i.i.d. random variables are equal to K times the mean and variance, respectively, of the individual random variables. This implies that the probability that campaign $(\hat{s}, c)$ is delivered at least $\tilde{m}_{\hat{s},c}$ impressions is

$$1-\Phi_N\left(\frac{\tilde{m}_{\hat{s},c} - K\hat{m}_{\hat{s},c}}{\sqrt{K\hat{m}_{\hat{s},c}}}\right) = 1-\Phi_N\left(\frac{-\sqrt{K\tilde{m}_{\hat{s},c}}z_\alpha}{\sqrt{K\hat{m}_{\hat{s},c}}}\right) = \Phi_N\left(\frac{\sqrt{K\tilde{m}_{\hat{s},c}}z_\alpha}{\sqrt{K\hat{m}_{\hat{s},c}}}\right) \approx \Phi_N\left(\sqrt{K}z_\alpha\right) \geq \alpha,$$

where we have used $\hat{m}_{\hat{s},c} = \frac{\tilde{m}_{\hat{s},c}}{K} + z_\alpha \sqrt{\frac{\tilde{m}_{\hat{s},c}}{K}} \approx \frac{\tilde{m}_{\hat{s},c}}{K}$, since $\frac{\tilde{m}_{\hat{s},c}}{K}$ is large.

Cost of Policy π_1: It is clear from the definition of π_1 that, for any period $s \in \mathbb{N}$, the expected cost incurred by π_1 (given the information available at the beginning of that period) is $G_1(\bar{\mathbf{m}}_\mathbf{s})$. For any $s \in \{K, K+1, \ldots\}$, we know that $\bar{m}_{s,c}$ is the realization of the random variable $\bar{M}_{s,c} := \sum_{\hat{s}=s-K+1}^{s} \hat{M}_{\hat{s},c}$, where $\hat{M}_{\hat{s},c} := \frac{\tilde{M}_{\hat{s},c}}{K} + z_\alpha \sqrt{\frac{\tilde{M}_{\hat{s},c}}{K}}$. By assumption, for every $c \in \mathcal{C}$, the set of random variables $\{\tilde{M}_{\hat{s},c} : \hat{s} \in \mathbb{N}\}$ is i.i.d. which implies that $\{\bar{M}_{s,c} : s \geq K\}$ is also i.i.d.. Let us use $\bar{M}_c$ to denote the sum of K i.i.d. samples of the random variable $\hat{M}_{1,c}$; then, $\bar{M}_{s,c} \sim \bar{M}_c$ for every $s \geq K$. Thus, the expected cost incurred in any period $s \geq K$, where the expectation is taken at the beginning of period 1, is $\mathbb{E}[G_1(\bar{\mathbf{M}})]$, where $\bar{\mathbf{M}}$ is the vector $(\bar{M}_c : c \in \mathcal{C})$. Therefore, the long run average expected cost per period is $\mathbb{E}[G_1(\bar{\mathbf{M}})]$. Combining this with the lower bound we derived on the optimal cost per period in Problem $\mathcal{P}_{Dyn}(\alpha)$, we obtain the following result.

Theorem 3. *The ratio of the expected cost per period incurred by π_1 to the lower bound on the optimal expected cost per period is*

$$\frac{\mathbb{E}[G_1(\bar{\mathbf{M}})]}{G_1(\alpha\mathbf{M})}. \tag{3.12}$$

Assessment of the Performance Guarantee in Theorem 3: We now examine the quality of our performance guarantee, i.e., (3.12), on the following realistic setup: Each time period represents one day, which consists of 20 million time slots. Each campaign has a time duration of 30 days (periods); thus $K = 30$. We consider five locations from the Boston area (locations 1, 2, 3, 4, and 5, corresponding to zip-codes 02110, 02114, 02116, 02118, and 02119, respectively) and estimate the expected-cost function $f_l(x)$ (of ensuring a win-probability of x at location l) in these five locations using Logistic-Regression (see the Appendix) on data obtained from Cidewalk.com. Figure 3.1 plots these functions for each of the five locations; consistent with our assumption in Section 2.1.1, these functions are indeed strictly increasing and strictly convex. For location $l \in \{1,\ 2,\ 3,\ 4,\ 5\}$, the expected-cost function is

$$f_l(x) = \frac{x}{\beta_1^l}\left[\ln\frac{x(1-x_0^l)+x_0^l}{1-x(1-x_0^l)-x_0^l} - \beta_0^l\right]$$

where $(\beta_0^1, \beta_0^2, \beta_0^3, \beta_0^4, \beta_0^5) = (-2.281, -2.192, -1.905, -1.936, -2.193)$, $(\beta_1^1, \beta_1^2, \beta_1^3, \beta_1^4, \beta_1^5) = (0.705, 1.042, 0.876, 0.767, 0.976)$ and $x_0^l = \frac{e^{\beta_0^l}}{1+e^{\beta_0^l}}$. We consider two values of the total number of locations L: 4 and 5. Each campaign is interested in a maximum of two locations; thus, when $L = 5$, there are a total of 15 possible sets of locations in which a campaign is interested (resp., 10 possible sets of locations when $L = 4$). Since the set of locations in which a campaign is interested defines its type, we have a total of $C = 15$ campaign-types when $L = 5$ (resp., $C = 10$ campaign-types when $L = 4$). In each period, a campaign of each type arrives with probability r; we consider the following five values of

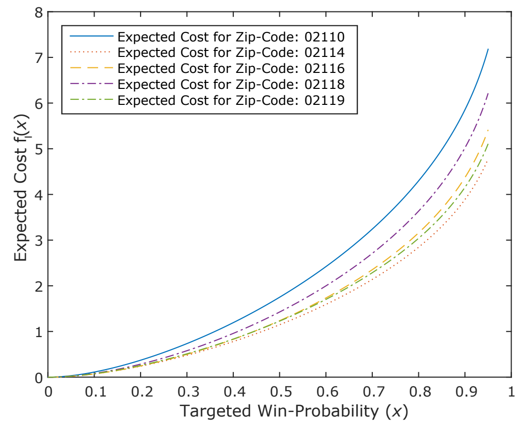

Figure 3.1. The expected cost $f_l(x)$ of ensuring a win-probability of x at five locations in the Boston area.

r: $\{0.1,\ 0.2,\ 0.3,\ 0.6,\ 0.9\}$. At each location, the probability that an impression arrives in a time slot is $q = 0.05$. The platform's probabilistic guarantee is $\alpha = 0.99$. For every $c \in \mathcal{C}$ and $s \in \mathcal{S}$, the requirement of campaign (s, c), conditional on that campaign arriving, is deterministic. This deterministic requirement takes one of the values in the set $\{10000, 50000, 150000, 250000\}$. In the following discussion, we refer to this parameter as the *campaign-demand*.

Thus, we have a total of $2 \times 5 \times 4 = 40$ different configurations of the parameters. The average (resp., maximum, minimum) value of the guarantee is 1.05 (resp., 1.09, 1.02). Thus, the cost of the rolling-horizon policy π_1 is at most 9% higher than the lower bound on the optimal cost which we developed in Section 3.1. The performance guarantee of our policy improves as the campaign-arrival probability parameter r increases and as the campaign-demand parameter increases. An increase in either parameter increases the "demand on the system" in the sense that the platform is required to procure a greater number of impressions. This leads to two effects: (1) The safety cushion (the square root term in the expression for $\hat{M}_{s,c}$) relative to the average demand per campaign, M_c (which is r times the campaign-demand), becomes smaller. (2) At the same time, the lower bound $G_1(\alpha\mathbf{M})$ is a convex increasing function of M_c. Together, these two effects cause the ratio in Theorem 3 to decrease. Moreover, an increase in r leads to an additional effect which also causes the ratio to decrease: (3) The number of impressions required by all type c campaigns which arrive over a certain number of periods is binomially distributed with mean proportional to r and variance proportional to $r(1-r)$. Thus, the coefficient of variation is proportional to $\sqrt{(1-r)/r}$, which decreases as r increases, thus making the problem closer to a deterministic problem.

The performance guarantee in Theorem 3 can be bounded above by an expression similar to the guarantee we obtained for the static model in Theorem 2. Recall from Section 2.1 that we defined $\psi_{i,l}(x) = \frac{xf'_{i,l}(x)}{f_{i,l}(x)}$. Let the support of the random variable $\tilde{M}_{\hat{s},c}$ be $[0,\ M_c^{max}]$. Let $\gamma^r_{max} = \max_{c \in \mathcal{C}} \left[\frac{M_c^{max} + z_\alpha \sqrt{KM_c^{max}}}{\alpha M_c} \right]$. As in our analysis in Section 2.1, since $\alpha \in [0.95,\ 0.99]$ and M_c^{max} is of the order of 100,000, we have $\gamma^r_{max} x_{max} \ll 1$. We proceed under the assumption that $\gamma^r_{max} x_{max} < 1$, and define $\overline{\psi}_r = \max\limits_{x \in [0, \gamma^r_{max} x_{max}]} \max\limits_{i \in \mathcal{I}, l \in \mathcal{L}} \psi_{i,l}(x)$. Similar to our assumption in Section 2.1.2, we assume here that $\overline{\psi}_r \leq \frac{\gamma^r_{max}}{\gamma^r_{max} - 1}$; again, it is easy to verify that this assumption holds comfortably in practice.

Corollary 3. *The ratio in Theorem 3 satisfies* $\frac{\mathbb{E}[G(\bar{\mathbf{M}})]}{G(\alpha \mathbf{M})} \leq \frac{1}{1 - \frac{(\gamma^r_{max} - 1)}{\gamma^r_{max}} \overline{\psi}_r}$.

3.3 Higher Values of the Rolling-Horizon Length

The performance bounds we derived in Theorem 3 and Corollary 3 were for a rolling-horizon length of one period. The rolling-horizon policy can be naturally extended for larger horizon lengths. The formal definition of this extension involves introducing cumbersome notation; we therefore provide a high-level description here. Let τ denote the length of the rolling horizon. At the beginning of, say, period s, the platform solves a τ-period problem (corresponding to periods $s, s+1, \ldots, s+\tau-1$), in which the demand of each active campaign is prorated to the horizon of τ periods. For example, if a campaign has a demand of p impressions over a remaining time duration of $q \geq \tau$ periods, then its prorated amount over τ time periods is $\frac{p}{q} \times \tau$. Future campaigns, expected to arrive within the horizon of τ periods, are prorated in a similar manner using their expected demand. After solving the τ-period problem, the platform implements only the decisions (i.e., the win and allocation probabilities) for the current period, i.e., period s. Thereafter, at the beginning of period $s+1$, the platform again solves a τ-period (corresponding to periods $s+1$, $s+2$, …, $s+\tau$) and implements the decisions for period $s+1$. The process continues.

While the performance guarantee of the rolling-horizon policy for $\tau = 1$ is quite attractive, it is instructive to examine its performance for higher values of τ. To assess this, we use the following setup: The total number of locations $L = 4$. Each campaign is interested in a maximum of two locations. Since the type of campaign is defined by the set of locations it is interested in, there are a total of 10 types of campaigns; thus, $C = 10$. The campaign arrival probability for each type of campaign is $r = 0.6$. For each of the four locations, the probability that an impression will arrive from that location in a time slot is $q = 0.05$. The platform's probabilistic guarantee $\alpha = 0.99$. We consider a time horizon of 100 time periods (days). Each campaign has a requirement of 100,000 impressions over a time duration of 10 periods; thus campaign-demand $= 100,000,\ K = 10$. At the beginning of each period, the platform solves a τ-period problem; we consider the following four values of τ: $\{1,\ 2,\ 3,\ 4\}$. Future campaigns are considered with their expected prorated demand within the τ-period horizon.

Figure 3.2 shows the change in the percentage gap between the average per-period cost (over the 50 simulations) of the rolling-horizon policy and the lower bound on the optimal per-period cost (derived in Section 3.1), for each choice of the rolling-horizon length τ. While the policy performs better for higher values of τ, the marginal improvement diminishes with τ.

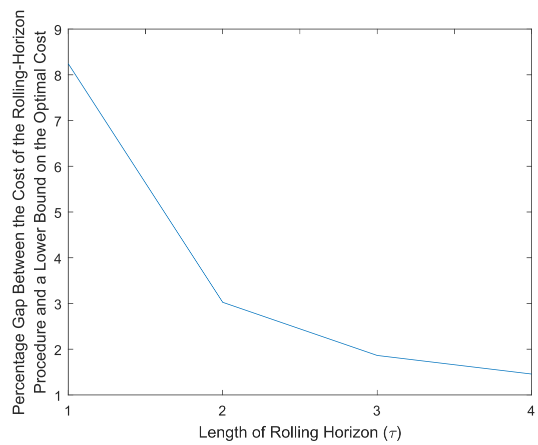

Figure 3.2. Percentage Gap Between the Cost of the Rolling-Horizon Procedure and the Lower Bound on the Optimal Cost Developed in Section 3.1.

3.4 Discussion and Future Research Directions

The procurement problem we addressed in this paper has the following salient features: (a) The supply is uncertain, (b) The supply cannot be inventoried, and (c) There are demand-side commitments which have to be met. Our approach can potentially be applied to a variety of other settings as well; we now illustrate with an example.

Third-Party Vendors at Google AdWords: Google's AdWords service allows advertisers to display their ads to users when they search on Google. Advertisers can bid to display their ads when the results of a keyword search are displayed to a user. Since all advertisers may not have the time and resources to build and maintain an AdWords account, Google runs a partner program through which approved third-party vendors (e.g., AOL, Adacado, AdSpeed) are allowed to bid on behalf of advertisers. A third-party vendor, who executes a campaign for an advertiser, faces the following problem: Given the set of search keywords of interest to the advertiser, the campaign's duration, and the campaign's budget, the vendor has to decide a bid for every impression that arises from the relevant keywords over the duration of the campaign, to maximize the advertiser's objective (e.g., maximize the number of impressions or clicks) without exceeding the budget.

As before, we divide the time horizon into time slots; let the duration of the campaign be T time slots. Let $\mathcal{W}$ be the set of keywords of interest to the advertiser and let Γ be the budget. Let q_w be the probability that an impression corresponding to keyword $w \in \mathcal{W}$ will arrive in a time slot. Let $\tilde{\Gamma}_t$ be the random variable representing the amount of money spent by the end of time slot t and let $b_{t,w}(\tilde{\Gamma}_{t-1})$ be the bid amount in period t for keyword w. Let $x_w(b)$ be the win-curve for keyword w; i.e., the win-probability of an impression corresponding to keyword w when the bid amount is b. Let α be the probability that the budget Γ will not be exceeded. The third-party vendor's optimization problem, when the advertiser's objective is, say, to maximize the number of impressions, is as follows:

$$\mathcal{P}_{AdWords}: \begin{cases} \max_{\mathbf{b}} \ \sum_{t=1}^{T} \sum_{w \in \mathcal{W}} \mathbb{E}\big[q_w x_{t,w}\big(b_{t,w}(\tilde{\Gamma}_{t-1})\big)\big] \\ \text{subject to: } \mathbb{P}\left[\tilde{\Gamma}_T > \Gamma\right] \leq 1 - \alpha, \end{cases}$$

where the evolution of $\tilde{\Gamma}_t$ is similar to that of $\vec{j}$ in the formulation of $\mathcal{P}_{static}(\vec{M}, \alpha)$ (Section 2.1). Our solution approach in this paper can be used to obtain a solution of problem $\mathcal{P}_{AdWords}$.

Our analysis in this paper did not consider the platform's *pricing* decisions for the various mobile advertising campaigns it offers. For given prices of the campaigns, our focus was on minimizing the procurement cost incurred by the platform in fulfilling the demand of these campaigns. One direction for future work is to study the platform's profit-maximization problem in which both the prices of the campaigns and the platform's procurement and allocation policies are decisions.

Our study focused on the operations of a "managed-service" platform that agrees to meet the demand-fulfillment constraints of the campaigns and obtains the required supply at least cost. As an area for future investigation, consider a problem arising in a "self-serve" mode, where the customer is completely in control, e.g., which apps to bid upon, how much to bid for each opportunity, when to stop a campaign, etc. In such a self-serve scenario, the mobile-promotion platform acts as a provider, with no operational decision-making responsibilities. Customers run their campaigns on a platform that faithfully executes whatever procurement strategy they choose to use. In this self-serve model, the problem we solved in this paper would shift to the customer and would presumably need to be solved on a much smaller scale. However, the analysis in the current study would be useful to guide end customers on how to optimize their mobile advertising campaigns.

CHAPTER 4

AD-BLOCKERS: A BLESSING OR A CURSE?

The stream of literature that is closest to our work is the one on *ad-avoidance*. The main focus of this literature is on the consequences of ad-avoidance for advertisers and consumers; see, e.g., Hann et al. 2008, Johnson 2013, and Goh et al. 2015. The modus operandi to avoid ads that have been examined here include, among others, avoiding TV ads by switching channels, avoiding tele-marketing ads by registering in Do-Not-Call lists, and ignoring marketing emails. There also exist studies that investigate the impact of ad-avoidance on the revenue of the facilitators of advertising content; e.g., TV broadcasters and print media (Stühmeier and Wenzel 2011). However, the ramifications of ad-avoidance on web-based advertising and the creation of effective response strategies for publishers is a relatively understudied subject. This is expected, since contracts for advertising on TV and telephone are long-term delivery contracts, where the facilitator usually gets paid regardless of whether or not the ad is actually seen by the consumer (in some cases, the payments are adjusted based on overall viewership statistics). Consequently, facilitators remain largely insulated from the ad-avoidance behavior of consumers. In contrast, a majority of present-day digital advertising content is cleared via real-time bidding for ad-space, with publishers often getting paid only when an ad is clicked upon by the consumer (i.e., pay-per-click payments). Therefore, in web-based advertising, ad-avoidance by consumers can significantly dent the revenue of publishers. With ad-blockers being the latest tool that consumers are increasingly adopting to avoid web ads, the question of how publishers should react intelligently to ad-blockers naturally gains importance. One key difference between ad-blocking on websites and ad-avoidance in traditional media (e.g., TV, radio) is the ability of facilitators to deliver ads: In contrast to traditional media, an ad cannot even be delivered on a website in the presence of an ad-blocker since it is blocked at its source. Further, relative to traditional media, web-based advertising presents a much richer setting to counter ad-avoidance since, here, the publisher has the ability to interact with each user *individually*; e.g., by asking a user to white-list the website. Accordingly, there is scope to develop novel "micro-level" strategies for publishers, investigate their structural properties, and assess their impact on revenue. Our work is an attempt to contribute in these directions. Two recent papers analyze settings that also involve the adoption of ad-blockers by users: Ray et al. (2017) analyze strategic interactions among content providers (websites), ad-blocking platforms, users, and advertisers. The content provider decides the quality of content and the ad-blocking platform decides the prices to charge users (for installation) and advertisers (for not blocking their ads). Given the quality of content provided by websites and the prices set by the platform, users and advertisers decide whether to adopt the platform. The paper derives useful insights on the optimal pricing structure for the platform, the adoption of the platform by users and advertisers, and the quality of content provided by websites. (Despotakis et al., 2017) analyze the impact of ad-blockers in a setting of two competing websites. Consistent with one of our results, this study also finds that the advent of ad-blockers can benefit the two websites; however, the analysis in this study does not consider the presence of network effects, which play a critical role in our analysis for deriving a website's optimal strategy in reacting to ad-block users.

The work that is perhaps closest to our study is (Hann et al., 2008). However, along with the above-mentioned finer granularity of decision-making in our context, there are several other important differences in the analysis. In (Hann et al., 2008), the value of the product being advertised does not depend on the number of users of the product. On the other hand, the network effect is at the core of our paper and plays a critical role in the website's decisions. Further, in (Hann et al., 2008), there are two consumer segments – *low-benefit* and *high-benefit* – that differ in the profitability they offer to the decision maker. In contrast, the consumer segments in our analysis are equally profitable from the website's viewpoint. It is in their individual sensitivities to ads (i.e., their respective ad-viewing costs) that these segments differ.

(Hann et al., 2008) find that when low-benefit consumers conceal themselves, advertising becomes cost effective and leads sellers to advertise more to the remaining users. One of our findings – actions by some consumers (i.e., ad-block users) to avoid advertisements affects other consumers – is consistent with the above result in (Hann et al., 2008). However, the force behind these two results is different. In (Hann et al., 2008), the externality arises from the advertising becoming cost effective. In our analysis, instead, the externality arises from (i) the ability of the website to discriminate users on the basis of their ad-viewing costs and (ii) the presence of the network effect. In practice, the extent of such externalities has been empirically established in Goh et al. (2015), in the context of consumers registering in Do-Not-Call lists to avoid marketing solicitations.

In our analysis, a website derives value from two components – an intrinsic component based on the website's content and another based on network effect (driven by the traffic/popularity of the website). This decomposition has been well-established in the literature, dating back at least to the work of Katz and Shapiro (1985). In their classic paper, Parker and Van Alstyne (2005) examine the reasons behind the practice of firms giving away free products. They argue that for firms that produce complementary goods, it may sometimes be profitable to provide a good for free. Intuitively, the cost incurred in providing a free good can be more than offset by the resulting increase in the demand for a premium good due to network effect. This result bears similarity to one of our results – in the ad-blocking context, allowing some ad-block users ad-free access increases the website's traffic and thereby its value (due to network effect), and in-turn leads to an increase in the number of regular (non-ad-block) users. The notion of network effect has been exploited in several other contexts; e.g., to examine the strategic interactions between service providers and their users (Nair et al., 2015), to measure the dependence between user-generated content and social ties in online social networks (Shriver et al., 2013), to measure the financial value of retaining and acquiring content contributors for a website which provides user-generated content (Zhang et al., 2012). Dou et al. (2013) examine how firms can engineer network externalities to their advantage and Kauffman et al. (2000) examine the impact of network externalities on the adoption of a network.

A website's decision to allow ad-block users to access the website without being subject to ads is akin to a firm providing a free version of its product or service; see, e.g., (Chellappa and Shivendu, 2010), Niculescu and Wu (2014), Lambrecht and Misra (2016). Typically, free versions offer a limited set of features, relative to the premium varieties. In our context, however, ad-blockers (if allowed to access the website) and regular users experience the same content.

The beneficial discrimination of the user population is an idea that our work shares with the versioning literature; see, e.g., Varian 1997, Chellappa and Shivendu 2005, Bhargava and Choudhary 2008, Wu and Chen 2008, Lahiri and Dey 2013, Wei and Nault 2014. There are significant differences, however, in our setting relative to what is typical in a traditional versioning problem, and also in how discrimination is operationalized. We now discuss some of these differences. Note that, in our context, all the users of the focal website have access to the same content – the basis for discrimination is the disutility they incur due to advertisements. In other words, all the users experience the *same consumption utility* (the equilibrium value of the website) but the two subgroups of users (namely, white-listers and regulars) incur *different disutilities* from ads. This is in contrast with the typical versioning setting, where price is the discriminating tool and a "premium" version usually offers a higher consumption utility than the "standard" version; e.g., the former typically offers more features than the latter. Further, in our setting, the absence of price as a discriminatory tool generates an interesting contrast in the following sense: In a typical versioning problem, the firm's revenue is determined by the consumption utility the product offers to the consumer (the price paid by the consumer equals the revenue the firm makes). In our setting, since users do not pay anything, it is the disutility imposed on them – and the extent of it – that indirectly generates revenue for the website (higher ad-intensities result in higher advertising revenue). Also, the notions of "high-type" and "low-type" consumers, which are common in the versioning literature, are not apparent in our setting. For instance, it is typical to find a high-type consumer incur more cost (since she buys the premium version, which is priced higher). On the other hand, in our context, a regular visitor is shown a relatively higher ad-intensity but his ad-sensitivity is relatively lower – thus, the cost (disutility) he incurs is not directly comparable to that for an ad-block user.

The discrimination of users based on their ad-viewing costs further gives rise to interesting operational possibilities such as selective gating (Section 4.5), where a subset of ad-block users is given ad-free access to exploit network effects (in the versioning setting, this would correspond to distributing a premium version free to a subset of consumers). One can also combine non-monetary and monetary levers to improve our capability to discriminate. For instance, in Section 4.6.2, we study a setting in which the website offers potential users the option of paid-subscription to access ad-free content in addition to the option of viewing ad-supported free content by white-listing the website. Thus, via the use of different ad-intensities, selective gating, and the subscription option, a website can have four types of users: (i) ad-block users who get ad-free access, (ii) white-listers (who get an ad-light experience), (iii) regular users (who get a relatively ad-heavy experience), and (iv) subscribers, who are given (paid) ad-free access.

A website that survives on advertising revenue can afford to let some ad-block users free-ride by enjoying ad-free content if it can generate sufficient revenue from non-ad-block users. In this sense, our work is related to the free-riding phenomenon. Shin (2007) analyzes the free-riding behavior of a discount retailer that does not provide pre-sale customer service (say, helping customers choose the right product) and therefore incurs lower operating costs, leading to lower prices. The author finds that when customers are heterogeneous in terms of their opportunity cost of visiting the discount retailer after receiving pre-sale customer service from another retailer that provides such a service, then the latter can benefit due to the free-riding behavior of the former. In our context too, the heterogeneity of users drives the increase in revenue of the website in the post-ad-block world. Asvanund et al. (2004) empirically examine the impact of free-riding by users in a P2P network and find that the extent of free-riding increases as the size of the network increases; Johar et al. (2011) analyze congestion in P2P networks and provide a mechanism to induce socially-optimal sharing.

The advent of ad-blockers has triggered an arms race between publishers and firms that develop ad-block software: developers continue to devise new techniques to evade the detection of their ad-blockers while publishers react by working out ways to detect an ad-blocker's presence. This has generated a significant amount of work in the Computer Science literature; for instance, new ad-blocking techniques (see, e.g., Krammer 2008, Storey et al. 2017), legal implications and impact on privacy (see, e.g., Krammer 2008, Vallade 2008, Gervais et al. 2016, and Garimella et al. 2017), and anti ad-blocking reactions and the impact of ad-blockers on the targeting ability of publishers (see, e.g., Johnson 2013 and Nithyanand et al. 2016).

4.1 Model

We consider a publisher (a website) whose entire revenue comes from advertisements and whose potential user-population consists of two segments: *ad-blockers* (potential users of the website, who use some ad-block software) and *regulars* (potential users of the website, who do not use any ad-block software).[1] We denote the size of the potential user-population by N, the fraction corresponding to ad-blockers by B, and the fraction corresponding to regulars by $\bar{B} := 1 - B$. We assume that the content-creation cost for this website is sunk, and the cost of serving that content to a user is zero. We consider the following two decisions for this publisher (see Figure 4.1 for a real-world illustration of these decisions):

- **Gating:** The publisher could allow ad-blockers to access the website for free (i.e., without showing any ads) or require them to white-list this website in order to access it. We use $I_G \in \{0, 1\}$ to denote this gating decision, with $I_G = 0$ denoting ad-free access and $I_G = 1$ the white-listing requirement.

- **Ad-Intensity:** If $I_G = 1$, the publisher decides two ad-intensities, a_r and a_b, respectively, for regulars and ad-blockers who white-list (we refer to these as *white-listers*). If $I_G = 0$, the publisher decides an ad-intensity a_r for regulars; since ad-blockers are allowed to access the website without seeing ads, $a_b = 0$ in this case. Driven by practical considerations (e.g., the minimum level necessary to grab user attention and/or the minimum amount dictated by the need to complete ad-campaigns in a timely manner), if ads are shown, the publisher is required to maintain a minimum ad-intensity[2] of $a_{\min} > 0$. Thus,

$$a_r, a_b \in \{0\} \cup [a_{\min}, \infty); \text{ moreover, } a_b = 0 \text{ if } I_G = 0.$$

While users do occasionally find ads useful, they are on *average* (based on past experience) typically perceived as a nuisance. Accordingly, we assume that users have a negative average perception of ads; in other words, users incur non-negative *ad-viewing costs*.

[1]For brevity, we refer to the publisher as "she" and a potential user as "he".

[2]The results for the case when there is no positive lower bound on the ad-intensities (i.e., $a_{\min} = 0$) are summarized later in Remark 2 (Section 4.3).

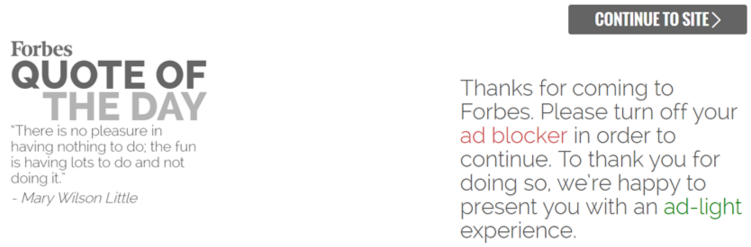

Figure 4.1. The message shown by Forbes.com to an ad-block user. The message highlights the following two actions: (i) The user is not allowed access to the website unless she disables the ad-blocker, and (ii) Upon white-listing, the user will be offered an ad-light experience.

We model heterogenous ad-viewing costs for potential users. For a randomly picked regular (ad-blocker) who is shown an ad-intensity of 1, the ad-viewing cost is a random variable denoted by $\tilde{c}_r \sim U[0,\ C_r]$ ($\tilde{c}_b \sim U[0,\ C_b]$). Therefore, when a regular (ad-blocker) is shown an ad-intensity of a_r (a_b), he incurs an ad-viewing cost of $a_r \tilde{c}_r$ ($a_b \tilde{c}_b$). Thus, the two segments of potential users differ in their distributions of ad-viewing costs for any given ad-intensity. We also refer to $\tilde{c}_b$ and $\tilde{c}_r$ as the *ad-sensitivity* of an ad-blocker and a regular user, respectively. We assume that $C_b \geq C_r$ to reflect the idea that ad-blockers are likely to perceive ads as being a greater nuisance than what regulars perceive. For every user visiting the website, the publisher makes a revenue of $r \cdot a$ if the ad-intensity is a; thus, r is the ad-revenue-per-user from an ad-intensity of 1.[3]

Let v denote the gross value that users obtain by accessing the website (excluding ad-viewing costs). This value, v, has two components: One component is the intrinsic value of the website generated through its content, and the other component is the network value of the website generated through the usage of the website. This decomposition

[3] Our analysis and results continue to hold for either of the two most popular payment formats in use today: pay-per-click and pay-per-impression. The revenue r can be interpreted as revenue per impression in the case of the pay-per-impression format or the expected revenue in the case of the pay-per-click format (click probability multiplied by per-click revenue).

as well as our assumption of uniformly distributed ad-viewing costs are consistent with the literature (see, e.g., Katz and Shapiro 1985). When users beneficially communicate and exchange information among themselves, the value of the website increases with the number of users. For instance, platforms such as Uber operate under very high two-sided network effects – an increase in the number of drivers increases the value of the platform for riders. Likewise, an increase in the number of riders increases the attractiveness of the platform for drivers. Platforms such as eBay and Amazon derive value from the interactions among their participants – more the number of participants, more is this value. An increase in the user base of a website increases its popularity and can, in turn, lead to an increase in the links (on other websites) that point to that website – this can improve its PageRank on Google search. Social media can also fuel network effects. For example, for a website such as Forbes.com, network value can be created by an increase in readership due to traffic enabled via social media by the sharing of content between existing users. Let n represent the number of users of this website out of the total population of N potential users. We assume the following functional form for v:

$$v = V\left[\bar{\alpha} + \alpha\frac{n}{N}\right], \tag{4.1}$$

where $\alpha \in [0,1]$ and $\bar{\alpha} := 1-\alpha$. Here, V denotes the maximum value users can obtain from the website. The fraction n/N denotes the extent of market penetration by this website. The constant $\bar{\alpha}$ denotes the fraction of the value generated from the website's content while $\alpha n/N$ denotes the corresponding fraction generated due to network effects. We therefore refer to α as the network-effect parameter.[4]

[4] The functional form $v = \overline{V}\left[\lambda + \gamma\frac{n}{N}\right]$ with $\lambda, \gamma \geq 0$, where the intrinsic value (represented by λ) and the network value (represented by γ) are independent, can be transformed to the form in (4.1) by letting $V = \overline{V}(\lambda+\gamma)$ and $\alpha = \frac{\gamma}{\lambda+\gamma}$. Here, the value of α can be interpreted as the *relative magnitude* of the network effect.

Let u_0 represent the net utility of an outside option for the potential users. The outside option represents some other website or source where similar content can be accessed (if no such source exists, then $u_0 = 0$). We assume that $u_0 \leq V$, failing which no potential user will be interested in visiting this website. Therefore, a regular with an ad-sensitivity of $\tilde{c}_r$ becomes a user of this website if $v - a_r\tilde{c}_r \geq u_0$. Similarly, an ad-blocker with an ad-sensitivity of $\tilde{c}_b$ becomes a user of this website if the website allows ad-free access to ad-blockers (i.e., does not gate them) and $v \geq u_0$ (in this case, $\tilde{c}_b$ is inconsequential); if the website requires ad-blockers to white-list (i.e., gates them), then this ad-blocker white-lists and becomes a user if $v - a_b\tilde{c}_b \geq u_0$. Thus, the publisher faces the following trade-off. On the one hand, an increase in the ad-intensity leads to a higher revenue from each user. On the other hand, this increase reduces the net utility that potential users obtain from this website and, therefore, reduces the number of users which, in turn, reduces revenue.

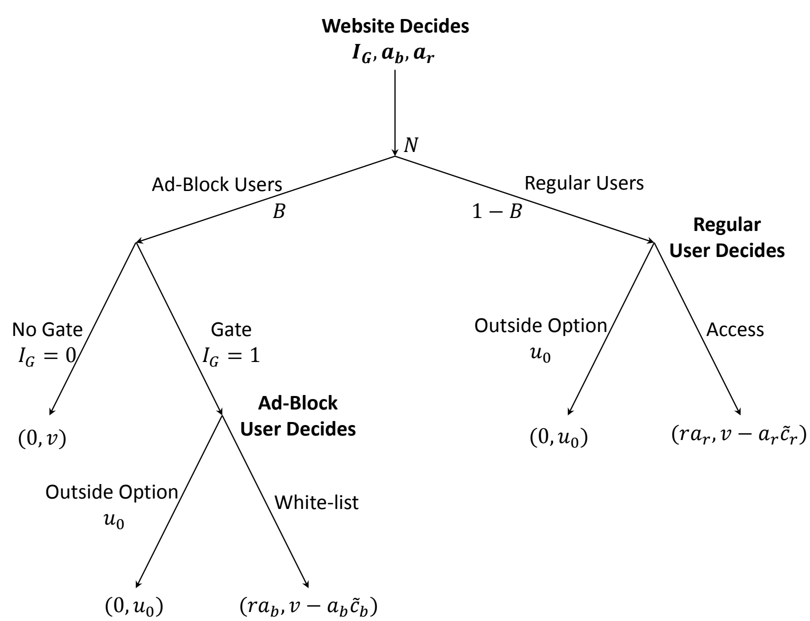

Figure 4.2. The sequence of events of the game: First the website decides $I_G,\ a_b,$ and a_r. Then, the ad-blockers decide whether or not to white-list the website, and the regulars decide whether or not to access the website. The quantities in brackets represent the payoffs received by the website and the user, respectively.

Table 4.1. The main notation for our analysis

Notation	Description
v	Equilibrium value of the website.
V	Maximum possible equilibrium value of the website.
n	Equilibrium traffic of the website.
N	Size of the potential user-population of the website.
α	Strength of the network effect. $\bar{\alpha} := 1 - \alpha$.
u_0	Net utility of the outside option.
B	Fraction of ad-block users among all potential users. $\bar{B} := 1 - B$.
I_G	Indicator variable representing the gating decision: Equals 1 if the website decides to gate; 0 otherwise.
a_b	Ad-intensity for white-listers.
a_r	Ad-intensity for regulars.
$a_{\min}$	Minimum possible (non-zero) ad-intensity.
C_b	Maximum ad-viewing cost among ad-block users for an ad-intensity of 1.
C_r	Maximum ad-viewing cost among regular users for an ad-intensity of 1.
r	Ad-revenue per user from an ad-intensity of 1.

Figure 4.2 depicts the sequence of events of the game. The website first decides I_G, a_b, and a_r. After the website's decisions, all the N potential users decide whether to pick the focal website (i.e., become a user) or pick the outside option, based on the comparison of their net utilities from these two sources as explained above. In Figure 4.2, the two quantities in brackets represent, respectively, the website's revenue from a potential user and that potential user's net utility. For every regular user, the website earns a revenue of ra_r. If $I_G = 0$ (that is, the website does not gate ad-blockers), the website makes no revenues from ad-blockers. If $I_G = 1$, the website makes a revenue of ra_b from every ad-blocker. Table 4.1 summarizes our main notation. Figure 4.3 is a graphical representation of the fraction of regulars who become users and the fraction of ad-blockers who become users.

Rational Expectations Equilibrium: Note that n, the number of users of the website, is the aggregate number of regulars and ad-blockers who *decide* to become users. These

Figure 4.3. An ad-block user with ad-viewing cost lower than $v - u_0$ white-lists the website, where v is the value users obtain from the website and u_0 is the net utility of the outside option. Similarly, a regular with ad-viewing cost lower than $v - u_0$ accesses the website.

individual decisions are influenced by the value v and their individual ad-viewing costs. However, v itself depends on n. We follow the standard assumption of *rational expectations* (see, e.g., Sheffrin 1996). That is, all potential users have a belief on v, the value of the website, and make their decisions. These decisions, in turn, result in a value for v which is consistent with the belief of the users.

Website's Problem: Every choice of $I_G,\ a_b$, and a_r will lead to a corresponding (i) equilibrium value of v denoted by $v(I_G, a_b, a_r)$, (ii) equilibrium number of ad-blockers who become users, denoted by $n_b(I_G, a_b, a_r)$, (iii) equilibrium number of regulars who become users, denoted by $n_r(I_G, a_b, a_r)$, and (iv) equilibrium value of the total traffic n, denoted by $n(I_G, a_b, a_r) = n_b(I_G, a_b, a_r) + n_r(I_G, a_b, a_r)$. Thus, the equilibrium revenue for the website, denoted by $R(I_G, a_b, a_r)$, can be expressed as follows:

$$R(I_G, a_b, a_r) = r\big[n_b(I_G, a_b, a_r) \cdot a_b + n_r(I_G, a_b, a_r) \cdot a_r\big]. \tag{4.2}$$

4.2 Analysis

In this section, we assess the impact of ad-block software on the website's decisions and its revenue. To this end, we solve for the optimal decisions of the website both *before* and *after* the advent of ad-block software, and also compare these decisions.

4.2.1 Optimal Decisions After the Advent of Ad-Blockers

The problem for the website after the advent of ad-block software is to maximize its equilibrium revenue; that is,

$$\max_{I_G, a_b, a_r} \quad R(I_G, a_b, a_r) \quad \text{s.t. } I_G \in \{0,1\} \text{ and } a_r, a_b \in \{0\} \cup [a_{\min}, \infty). \tag{P_{After}}$$

Let I_G^*, a_b^* and a_r^*, respectively, denote the optimal values of I_G, a_b and a_r, and let $v^* := v(I_G^*, a_b^*, a_r^*)$. We first note some important properties of the optimal solution (I_G^*, a_b^*, a_r^*).

- **Property A1: (Ad-Light Experience for White-Listers)** The optimal ad-intensities are such that $a_b^* \leq a_r^*$.

  - **Intuition:** We know that $C_b \geq C_r$, i.e., ad-blockers are more sensitive to ads than regulars. Therefore, the optimal ad-intensity for ad-blockers is less than that for regulars. This property is derived as part of the proof of Theorem 5 in the Appendix. □

- **Property A2: (Complete Market Penetration)** If $a_b^* > a_{\min}$, then $v^* - a_b^* C_b = u_0$ and $v^* - a_r^* C_r = u_0$; that is, all regulars and ad-blockers become users.

  - **Proof:** Since $a_b^* > a_{\min}$, we should have $I_G^* = 1$; moreover, from Property A1, we know that $a_r^* > a_{\min}$. Thus, Problem P_{After} simplifies to the unconstrained maximization of $R(1, a_b, a_r)$. The number of ad-blockers who become users and the number of regulars who become users are both proportional to the "value premium" $v - u_0$ and inversely proportional to the respective ad-intensities. That is, $n_b(1, a_b^*, a_r^*) = \frac{NB(v^* - u_0)}{a_b^* C_b}$ and $n_r(1, a_b^*, a_r^*) = \frac{N(1-B)(v^* - u_0)}{a_r^* C_r}$; therefore, $R(1, a_b^*, a_r^*) = n_b(1, a_b^*, a_r^*) r a_b^* + n_r(1, a_b^*, a_r^*) r a_r^*$. Thus, maximizing revenue is equivalent to maximizing v which, in turn, is equivalent to maximizing the total traffic, n. To maximize n, we need to convert all ad-blockers and regulars to users. The largest ad-intensities which will allow this are those that ensure that the regular with the highest ad-viewing cost and the ad-blocker with the highest ad-viewing cost are both indifferent between becoming a user and choosing the outside option. The statement of the property above ensures exactly that. □

- **Property A3: (Rational Belief on Value)** The optimal solution satisfies the following relations: $0 \leq v^* - u_0 \leq a_r^* C_r$ and, if $I_G = 1,\ v^* - u_0 \leq a_b^* C_b$.

  - **Proof:** (i) If $v^* < u_0$, then no potential user visits the website; thus, the revenue obtained is zero and cannot be optimal – thus, we should have $0 \leq v^* - u_0$. (ii) If $v^* - u_0 > a_r^* C_r$, then all regulars become users of the website because the net utility obtained by that regular with the highest ad-viewing cost is $v^* - a_r^* C_r$ which *strictly exceeds* u_0, the value of the outside option. Thus, the website could increase a_r^* slightly and still ensure that all regulars continue to be users of the website, while the net utilities obtained by ad-blockers remain unchanged (and thus, the number of ad-blockers who visit the website also remains unchanged). Thus, this increase in a_r^* will lead to a revenue increase. This shows that, at the optimal solution, $v^* - u_0 \leq a_r^* C_r$. An identical argument shows that, at the optimal solution, $v^* - u_0 \leq a_b^* C_b$. □

- **Property A4: (Equilibrium Traffic)** The equilibrium number of regulars and the equilibrium number of ad-blockers who become users are given by the following expressions:

$$
\begin{aligned}
n_r(I_G^*, a_b^*, a_r^*) &= N\bar{B}(v^* - u_0)/(a_r^* C_r), \\
n_b(I_G^*, a_b^*, a_r^*) &= NB(v^* - u_0)/(a_b^* C_b), \text{ if } I_G^* = 1, \text{ and} \\
n_b(I_G^*, a_b^*, a_r^*) &= NB, \text{ if } I_G^* = 0.
\end{aligned}
$$

- **Proof:** This result follows directly from Property A3 and our model of uniformly distributed ad-viewing costs for regulars and ad-blockers. □

The revenue function $R(I_G, a_b, a_r)$ can now be computed using the expressions in Property A4. Theorem 4 below states the optimal solution to problem P_{After}; a proof of optimality is provided in the Appendix. To express the optimal solution, we define the following four parametric constants:

$$
u_m = V - \frac{V\alpha B(BC_r + \bar{B}C_b)a_{\min}}{\bar{B}(BV\alpha - a_{\min}C_b) + (BC_r + \bar{B}C_b)a_{\min}},
$$

$$
a_g = \max\left\{a_{\min},\ \frac{\big[V(1-\alpha B) - u_0\big]C_b a_{\min}}{C_r(C_b a_{\min} - V\alpha B)}\right\},
$$

$$
\widehat{C} = \frac{V - u_0}{a_{\min}}, \quad \text{and}
$$

$$
u_h = V\bar{\alpha} - \frac{V\alpha B K_2}{K_3 - K_2},
$$

where $K_2 = \frac{\bar{B}a_{\min}}{a_{\min}C_r - V\alpha\bar{B}}$ and $K_3 = \frac{(BC_r + \bar{B}C_b)a_{\min}}{a_{\min}C_bC_r - V\alpha(BC_r + \bar{B}C_b)}$.

For a chosen set of decisions (I_G, a_b, a_r), the value premium offered by the website is $\phi(I_G, a_b, a_r) = v(I_G, a_b, a_r) - u_0$. By expressing the equilibrium value v in terms of the decisions (see (A.20) in the Appendix), we obtain

$$
\phi(I_G, a_b, a_r) = \frac{V\bar{\alpha} + V\alpha(1 - I_G)B - u_0}{1 - V\alpha\left(\frac{I_G B}{C_b a_b} + \frac{1-B}{C_r a_r}\right)}.
$$

The condition $\phi(I_G, a_b, a_r) \geq 0$, which appears in the result below, can be naturally interpreted as the "survival condition" for the website, in the sense that the website does not receive any traffic if this condition fails.

Theorem 4. *An optimal solution of problem P_{After} is as follows:*

- *If $C_b,\ C_r < \widehat{C}$, then*

$$(I_G^*,\ a_b^*,\ a_r^*) = \left(1,\ \frac{V-u_0}{C_b},\ \frac{V-u_0}{C_r}\right).$$

- *If $C_b,\ C_r > \widehat{C}$, then*

$$(I_G^*,\ a_b^*,\ a_r^*) = \begin{cases} (0,\ a_{\min},\ a_{\min}), & \textit{if } u_0 > u_h,\ \phi(0,\ a_{\min},\ a_{\min}) \geq 0, \\ (1,\ a_{\min},\ a_{\min}), & \textit{if } u_0 \leq u_h,\ \phi(1,\ a_{\min},\ a_{\min}) \geq 0, \\ (0,\ 0,\ 0), & \textit{otherwise}. \end{cases}$$

- *If $C_r \leq \widehat{C} \leq C_b$, then*

$$(I_G^*,\ a_b^*,\ a_r^*) = \begin{cases} (0,\ a_{\min},\ \frac{V-u_0}{C_r}), & \textit{if } u_0 > u_m,\ \phi(0,\ a_{\min},\ \frac{V-u_0}{C_r}) \geq 0, \\ (1,\ a_{\min},\ a_g) & \textit{if } u_0 \leq u_m,\ \phi(1,\ a_{\min},\ a_g) \geq 0, \\ (0,\ 0,\ 0), & \textit{otherwise}. \end{cases}$$

The three cases in this result signify three types of websites, categorized on the basis of the ad-viewing costs (or, equivalently, ad-sensitivities) of their user-population. We elaborate below:

(a) Low C_b and Low C_r $\left(C_b,\ C_r < \widehat{C}\right)$: This represents the class of websites whose target audience is not very sensitive to ads; for instance, websites that offer discounted deals. Since potential users of such websites are likely to white-list when required to do so, it is optimal to always gate ad-blockers. Moreover, the higher ad-tolerance of the users implies that the website can set relatively high ad-intensities (as compared to the two cases below) for both regulars and ad-blockers.

(b) High C_b and High C_r $\left(C_b,\ C_r\ >\ \widehat{C}\right)$: This represents websites which cater to an ad-sensitive population – for example, senior business executives. An ad-sensitive audience necessitates a cautious gating decision: Gate ad-blockers only if the outside option is not very attractive. Understandably, the ad-intensity is at its lowest for both regulars and ad-blockers.

(c) High C_b and Low C_r $\left(C_r \leq \widehat{C} \leq C_b\right)$: In each of the previous two cases, ad-blockers and regulars were similar in that both groups had low (high) ad-sensitivity. This third case represents websites that face a wider variety of ad-sensitivities in their potential users. Here, when the outside option is attractive (i.e., when $u_0 > u_m$), it is optimal not to gate ad-blockers. Further, since the ad-sensitivity of regulars is low, as in Case (a) above, the optimal ad-intensity for these users is the same as in that case. When the outside option is not attractive (i.e., when $u_0 \leq u_m$), it becomes optimal to gate ad-blockers. Under this possibility, since ad-blockers are highly ad-sensitive, their ad-intensity is the lowest possible. The gating of ad-blockers results in the website losing some of these potential users. To compensate, the website attracts regulars by offering them an ad-intensity which is lower than that in Case (a) (it is easy to see that $a_g \leq \frac{V-u_0}{C_r}$).

The closed-form expressions in Theorem 4 help us establish the monotonicity of the ad-intensities, for a given gating decision, with respect to an increase in the upper bound V on the value of the website. Surprisingly, as shown in Theorem 5 below, the optimal gating decision may *not* be monotone in the value of the website.

Theorem 5. (Behavior of Ad-Intensities and Gating Decisions): *(i) For a given gating decision I_G, the optimal ad-intensities $(\hat{a}_b(I_G), \hat{a}_r(I_G))$ increase with the maximum value V offered by the website. Thus, for a given gating decision, the website advertises more as its maximum value increases. (ii) The optimal gating decision (I_G^*) is, in general, not monotone in V.*

Interestingly, a website may decide to stop gating with an increase in V, under certain conditions. Consider, for instance, a website facing (i) a high network effect (i.e., α is high), (ii) a low ad-blocking rate (i.e., B is low), and (iii) a highly ad-sensitive target audience (i.e., C_b and C_r are high). As can been seen from Theorem 4, if the website gates, then a highly ad-sensitive user population makes it optimal to advertise at the minimum level both to regulars and to ad-blockers who white-list. Here, ad blockers contribute little to the revenue of the website, since they constitute a small fraction of the potential user population and, further, being ad-sensitive, most of them do not white-list. In this situation, an increase in V can make it optimal for the website to stop gating, driven by the following benefits: (i) Due to a strong network effect, allowing ad-free access to ad-blockers increases the value of the website (over and above the increase due to a higher V). (ii) In turn, this increase in value enables the website to attract more regulars and also advertise to them at a higher intensity. Note that a further increase in V may make it optimal for the website to start gating again: If the increase in V is substantial enough, then despite the loss of ad-blockers who refuse to white-list, the website can (a) further increase the advertising intensity for regulars and (b) generate revenue from ad-blockers who white-list.

Remark 1. *The special case of $\alpha = 0$ corresponds to the complete absence of the network effect. In this case, it is easy to verify that $u_h = u_m = V$ and therefore, from Theorem 4, the website's optimal strategy is to always gate ad-block users. More generally, this strategy is attractive for websites that experience a low network effect; e.g., forbes.com and businessinsider.com. On the other hand, a website such as facebook.com relies heavily on the network effect and would therefore refrain from always gating ad-block users.*

To be able to assess the impact of ad-block software on the ad-intensities and the revenue of the website, we next analyze the website's decisions before the advent of such software.

4.2.2 Optimal Ad-Intensity in the Pre-Ad-Block World

After the advent of ad-block software, the publisher is able to discriminate between the two segments of potential users, namely regulars and ad-blockers, who differ in their ad-viewing cost distributions. In the world prior to ad-block software, the publisher had no such tool to identify which segment a potential user belongs to and thus used a common ad-intensity for all potential users. Since it could be confusing to refer to the two segments as regulars and ad-blockers in the pre-ad-block world, we refer to these segments as *low-cost* and *high-cost*, respectively. The subscript r will refer to the low-cost segment and the subscript b will refer to the high-cost segment. Using the same notation as in the previous section, the pre-ad-blocking world corresponds to the website always gating; i.e., $I_G = 1$ and choosing a common ad-intensity $a_b = a_r = a$ (say). Thus, the website's revenue is $R(1, a, a)$ and we now have the following optimization problem, which we refer to as problem P_{Before}.

$$\max_{a} \quad R(1, a, a) \quad \text{s.t. } a \in \{0\} \cup [a_{\min}, \infty). \tag{P_{Before}}$$

Let a^* be the optimal ad-intensity in problem P_{Before}. Let $v^* := v(1, a^*, a^*)$ (resp., $n(1, a^*, a^*)$) denote the corresponding equilibrium value (resp., equilibrium traffic) of the website. As in Section 4.2.1, we first present important properties that an optimal solution of problem P_{Before} must satisfy and then use them to derive an optimal solution.

- **Property B1: (Ad Intensities)** Let $\hat{u}_0 = V\left[1 - \frac{\alpha B\bar{\rho}}{1-B}\right]$. If $a^* > a_{\min}$, then (i) if $u_0 \leq \hat{u}_0$, the optimal ad-intensity a^* satisfies $v^* - a^* C_r = u_0$, (ii) otherwise, $v^* - C_b a^* = u_0$.
  - **Intuition:** When the utility of the outside option is low, the website can hope to attract users even at a high ad-intensity. Driven by this, the optimal ad-intensity is such that the entire low-cost segment prefers the website to the outside option but only a fraction of the high-cost segment does so. When

the outside option is attractive, the website is more conservative and uses a relatively lower ad-intensity, which ensures that both the low-cost and high-cost segments become users.

This property is derived as part of the proof of Theorem 6 in the Appendix. □

- **Property B2: (Rational Belief on Value)** The optimal ad-intensity a^* and the equilibrium value v^* satisfy $0 \leq v^* - u_0 \leq a^* C_b$.

  - **Proof:** The proof of this property is similar to that of Property A3 in Section 4.2.1. □

- **Property B3: (Equilibrium Traffic)** The equilibrium number of users is

$$n(1, a^*, a^*) = \begin{cases} N(1-B)\frac{(v^*-u_0)}{a^* C_r} + NB\frac{(v^*-u_0)}{a^* C_b}, & \text{if } 0 \leq v^* - u_0 \leq a^* C_r, \\ \\ N(1-B) + NB\frac{(v^*-u_0)}{a^* C_b}, & \text{if } a^* C_r \leq v^* - u_0 \leq a^* C_b. \end{cases} \tag{4.3}$$

  - **Proof:** This result follows directly from Property B2 and our model of uniformly distributed ad-viewing costs for high-cost and low-cost users. □

The revenue function $R(1, a, a)$ can now be computed using the expression of $n(1, a^*, a^*)$ in (4.3). Theorem 6 below states the optimal solution to problem P_{Before}; a proof of optimality is provided in the Appendix. To express the optimal solution, we define the following four parametric constants:

$$\widehat{C}_r = \frac{(V - u_0 - V\alpha B)C_b}{a_{\min} C_b - V\alpha B},$$
$$u_b = V(1 - \alpha B) + \frac{M_2}{M_1},$$
$$\hat{a} = \frac{V(1 - \alpha\rho) - u_0}{C_r},$$
$$\hat{\hat{a}} = \frac{V - u_0}{C_b},$$

where $\rho = \frac{B(C_b - C_r)}{C_b}$, $\bar{\rho} = 1 - \rho$, $M_1 = \frac{\bar{\rho}}{C_r} + \frac{Ba_{\min}}{V\alpha B - C_b a_{\min}}$, and $M_2 = \frac{\bar{\rho}V\alpha - \bar{B}a_{\min}C_b}{C_b}$. It is easy to show that $\hat{a} \geq \hat{\hat{a}}$ and $\widehat{C} \geq \widehat{C}_r$.

As we did in Section 4.2.1, it is convenient to obtain the expression for the value premium, i.e., $v(a) - u_0$, offered by the website, for a chosen ad-intensity a. Using the expression for the equilibrium value v $\big($see (A.30) and (A.34) in the Appendix$\big)$, we have

$$v(a) - u_0 = \begin{cases} \phi_l(a) = \frac{aC_r(V\bar{\alpha} - u_0)}{aC_r - V\alpha\bar{\rho}}, & \text{if } C_b \geq \widehat{C},\ C_r \geq \widehat{C}_r, \\ \\ \phi_h(a) = \frac{aC_b[V(1-\alpha B) - u_0]}{aC_b - V\alpha B}, & \text{otherwise.} \end{cases}$$

Parallel to the meaning of the condition $\phi(I_G, a_b, a_r) \geq 0$ in the analysis of Problem P_{After} (Section 4.2.1), here the survival condition for the website (i.e., the website does not receive any traffic if this condition fails) is (i) $\phi_l(a) \geq 0$, for a highly ad-sensitive user population and (ii) $\phi_h(a) \geq 0$, otherwise. We use these conditions in the result below.

Theorem 6. *An optimal solution of problem P_{Before} is as follows:*

- *If $C_b,\ C_r < \widehat{C}$, then*

$$a^* = \begin{cases} \hat{a}, & \textit{if } u_0 \leq \hat{u}_0,\ \phi_h(\hat{a}) \geq 0, \\ \hat{\hat{a}}, & \textit{if } u_0 > \hat{u}_0,\ \phi_h(\hat{\hat{a}}) \geq 0, \\ 0, & \textit{otherwise.} \end{cases}$$

- *If $C_b \geq \widehat{C},\ C_r \geq \widehat{C}_r$, then*

$$a^* = \begin{cases} a_{\min}, & \textit{if } \phi_l(a_{\min}) \geq 0, \\ 0, & \textit{otherwise.} \end{cases}$$

- *If* $C_b \geq \widehat{C},\ C_r < \widehat{C}_r$*, then*

$$a^* = \begin{cases} \hat{a}, & \textit{if } u_0 \leq u_b,\ \phi_h(\hat{a}) \geq 0, \\ a_{\min}, & \textit{if } u_0 > u_b,\ \phi_h(a_{\min}) \geq 0, \\ 0, & \textit{otherwise}. \end{cases}$$

As in Theorem 4, the three cases in the above result represent three categories of websites based on the ad-sensitivities of their potential users.

(a) Low C_b **and Low** C_r $\left(C_b,\ C_r < \widehat{C}\right)$**:** When the outside option is low (high), the website uses a relatively high ad-intensity $\hat{a}$ (resp., low ad-intensity $\hat{\hat{a}}$). Due to the low ad-sensitivity (i.e., high ad-tolerance) of the user-population in this category, these intensities are higher than their corresponding values for the other two categories.

(b) High C_b **and High** C_r $\left(C_b \geq \widehat{C},\ C_r \geq \widehat{C}_r\right)$**:** Here, due to the high ad-sensitivity of the user-population, it is optimal to keep the ad-intensity at its lowest.

(c) High C_b **and Low** C_r $\left(C_b \geq \widehat{C},\ C_r < \widehat{C}_r\right)$**:** In this case, the website serves users with a wider variety of ad-sensitivities. If the outside option is not attractive, then the website advertises at a high ad-intensity ($\hat{a}$); otherwise, it chooses the the minimum ad-intensity.

Analogous to Theorem 5 (which was for Problem P_{After}), the following result establishes the monotonicity of the optimal ad-intensity in Problem P_{Before} with respect to the upper bound V on the value of the website. The proof of this result is similar to that of Theorem 5 and therefore not provided for brevity.

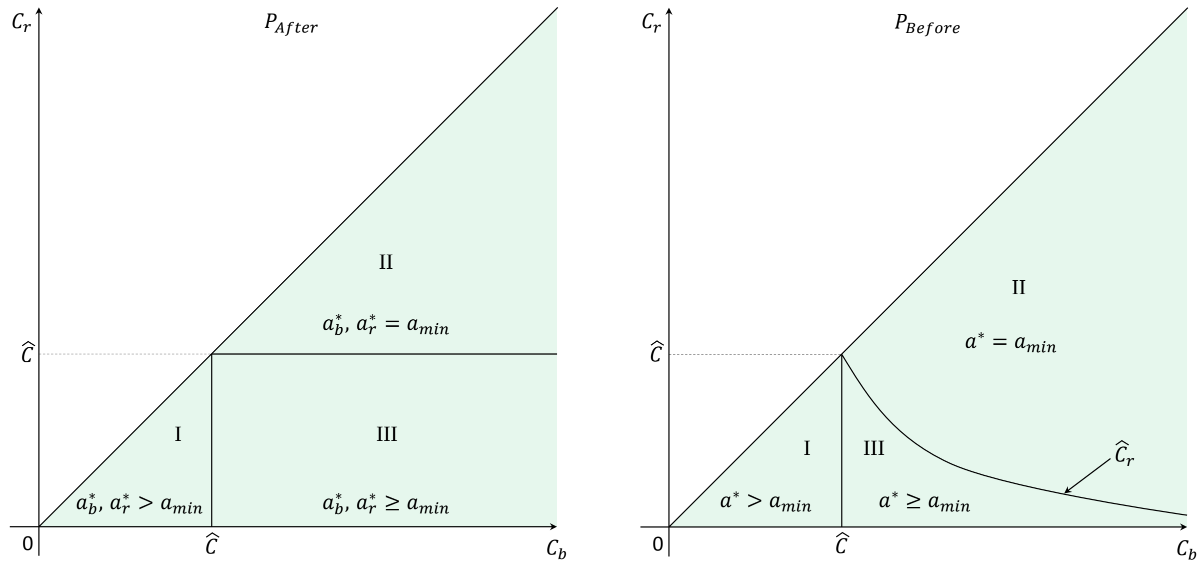

Figure 4.4. A pictorial comparison of the ad-intensities in the optimal solutions to problems P_{After} and P_{Before}.

Theorem 7. (Monotonicity of Ad-Intensity): *Keeping all other parameters fixed, as the maximum value V offered by the website increases, the optimal ad-intensity a^* increases. In other words, the website advertises more as its maximum value increases.*

Using the optimal solutions in Theorems 4 and 6, we now compare the ad-intensities before and after the advent of ad-block software.

4.2.3 Comparing Optimal Ad-Intensities in the Pre- and Post-Ad-Block Worlds

Figure 4.4 is a pictorial comparison of the ad-intensities in the optimal solutions to problems P_{After} and P_{Before}. As shown here, for both these problems, the (C_b, C_r)-space can be decomposed into three regions – I, II, and III – that broadly represent the websites that serve users with low, high, and moderate ad-sensitivities, respectively. In the pre-ad-block world, the websites belonging to Region II found it optimal to keep the ad-intensity at its minimum level; i.e., $a_{\min}$. In the post-ad-block world, the discriminatory power bestowed by ad-blockers enables the choice of a higher ad-intensity for regular users with

$C_b \geq \widehat{C},\ C_r \in [\widehat{C}_r,\ \widehat{C}]$. Thus, Region II of the pre-ad-block world shrinks in the post-ad-block world (or, equivalently Region III expands).

4.3 Impact of Ad-Block Software on Publishers

In the previous section, we obtained the optimal decisions of the publisher before and after the advent of ad-block software. In this section, we use the analysis thus far to assess the impact on (i) the revenue of the website and (ii) the ad-intensities for the two segments of the population.

Theorem 8. *Under the optimal decisions derived in Section 4.2, the revenue of the website increases after the advent of ad-block software. That is,*

$$R(I_G^*, a_b^*, a_r^*) \geq R(1, a^*, a^*).$$

The intuition behind the above result is as follows. By definition, a^* is the optimal ad-intensity of the website before the advent of ad-block software. Consider the following strategy of the website after the advent of ad-block software: $a_b = a_r = a^*,\ I_G = 1$. In following this strategy, the website continues to keep the same ad-intensity after the advent of ad-block software as before, and gates all ad-blockers. Since the ad-intensity is the same before and after the advent of ad-block software and the website gates all ad-blockers, any potential user who became a user in the pre-ad-block world also becomes a user in the post-ad-block world. Thus, the maximum revenue obtained by the publisher before the advent of ad-block software is a lower bound on the optimal revenue of the website after the advent of ad-block software. This finding is consistent with that of a recent empirical investigation at Forbes (DVorkin 2016) in which ad-blockers were gated and offered an ad-light experience while regulars were offered a higher ad-intensity. The main finding was that, as compared to the earlier setting where ad-blockers were not gated, Forbes was able to deliver 63 million additional ad impressions in two weeks to those users who agreed to whitelist, without significantly affecting the total traffic to the website.

Theorem 9. *Compared to the ad-intensity before the advent of ad-block software, the ad-intensity for the high-cost segment decreases and the ad-intensity for the low-cost segment increases after the advent of ad-block software. That is,*

$$a_b^* \leq a^* \leq a_r^*.$$

The intuition for the above theorem is this: Since the two segments of the population differ in terms of their ad-viewing cost distributions, the ability to discriminate between their ad-intensities, in the world with ad-block software, makes it optimal to more aggressively advertise to the low-cost segment and to less aggressively advertise to the high-cost segment, than before.

Remark 2. *(Unconstrained Ad-Intensities)* Both problems P_{Before} and P_{After} imposed a lower bound, namely $a_{\min}$, on the ad-intensities. In the absence of this lower bound, i.e., when $a_{\min} = 0$, the optimal solutions of these problems simplify significantly. Using the same notation as in our analysis above, the optimal solution of P_{After} when $a_{\min} = 0$ is

$$(I_G^*,\ a_b^*,\ a_r^*) = \left(1,\ \frac{V - u_0}{C_b},\ \frac{V - u_0}{C_r}\right).$$

and the optimal solution of P_{Before} is

$$a^* = \begin{cases} \hat{a}, & \text{if } u_0 < \hat{u}_0, \\ \hat{\hat{a}}, & \text{if } \hat{u}_0 \leq u_0 \leq V\bar{\alpha} + V\alpha\bar{B}, \\ 0, & \text{otherwise.} \end{cases}$$

The key results of our analysis continue to hold when $a_{\min} = 0$: the website's revenue increases, white-listers receive an ad-light experience, and regulars are offered a higher ad-intensity in the post-ad-block world. □

Thus far, our focus was on the impact of ad-block software on the website's revenue. We now examine the impact on *consumer surplus* and *social surplus*.

4.4 Welfare Analysis

We define consumer surplus as the total expected net utility over all the consumers; that is, the sum of the net utilities of all the consumers. Social surplus is defined as the sum of consumer surplus and the revenue of the website. The main takeaway from our analysis in this section is an argument in favor of a website providing niche content to ensure that both the website and the consumers of its content benefit. We now proceed with the analysis.

For simplicity of exposition, we focus on the case when the minimum ad-intensity threshold is low; specifically, $a_{\min} \leq \frac{V-u_0}{C_b}$. From our analysis of Problem P_{After} in Section 4.2 we know that, in this case, $a_b^* = \frac{V-u_0}{C_b},\ a_r^* = \frac{V-u_0}{C_r}$, $n = N$, and $v = V$. Let CS_{After} denote the consumer surplus after the advent of ad-block software. An ad-block user with an ad-sensitivity of $\tilde{c}_b$ receives a net utility of $V - a_b^* \tilde{c}_b$. Similarly, a regular user with an ad-sensitivity of $\tilde{c}_r$ receives a net utility of $V - a_r^* \tilde{c}_r$. Since $\tilde{c}_b \sim U[0,\ a_b^* C_b]$ and $\tilde{c}_r \sim U[0,\ a_r^* C_r]$, we have

$$CS_{After} = NB\left[V - \frac{a_b^* C_b}{2}\right] + N(1-B)\left[V - \frac{a_r^* C_r}{2}\right].$$

Using $a_r^* = \frac{V-u_0}{C_r}$ and $a_b^* = \frac{V-u_0}{C_b}$ we get

$$CS_{After} = \frac{N(V+u_0)}{2}. \tag{4.4}$$

Let CS_{Before} denote the consumer surplus before the advent of ad-block software. Using an analysis similar to that above, we have

$$CS_{Before} = \begin{cases} \frac{N\bar{\rho}(V - V\alpha\rho + u_0) + 2NB\rho u_0}{2}, & \text{if } u_0 \leq \hat{u}_0, \\ \\ \frac{N\bar{B}[2C_b V - C_r(V - u_0)] + C_b NB(V + u_0)}{2C_b}, & \text{otherwise.} \end{cases} \tag{4.5}$$

For more details on the expressions in (4.5), we refer the reader to the proof of the following result, which characterizes the increase in consumer surplus after the advent of ad-block software.

Theorem 10. *After the advent of ad-block software, consumer surplus increases if and only if* $u_0 \leq \hat{u}_0$.

The intuition behind this result is as follows. In the pre-ad-block world, the publisher – who had no tool to segment users based on their ad viewing costs – used a single ad-intensity for all potential users. Not surprisingly, this common ad-intensity is too high for some high-cost potential users, who therefore choose not to become users of the website and instead consume the low-utility outside option ($u_0 \leq \hat{u}_0$). However, after the advent of ad-block software, the website is able to discriminate by offering a lower ad-intensity to such high-cost users, thereby converting them into users of the website and leading to an increase in the total traffic of the website. This increases consumer surplus in two ways: First, all the new users (the users of the website after the advent of ad-block software but not before) obtain a higher net utility compared to the low utility from the outside option, which they were consuming earlier. Second, the increase in traffic increases the value of the website; therefore, the old users (the users of the website both before and after the advent of ad-block software) also obtain a higher net utility from the website. In this manner, consumer surplus increases after the advent of ad-block software when the outside option has low net utility.

When the outside option offers a high-enough net utility ($u_0 > \hat{u}_0$), then the ad-intensity in the world prior to ad-blockers is itself quite low (see Theorem 6). After the advent of ad-block software, while the website is able to decrease the ad-intensity for high-cost users, this decrease is not significant given that the ad-intensity was already low. On the other hand, the ad-intensity increases significantly for low-cost users – while the website is able to retain these users by carefully increasing their ad-intensity, their consumer surplus decreases due to this increase. Overall, the loss of consumer surplus due to the (significant) increase in the ad-intensity for low-cost users more than offsets the gain in consumer surplus due to the (small) decrease in the ad-intensity for high-cost users. As a consequence, the total consumer surplus decreases.

We now examine the social surplus. Let W_{Before} and W_{After} (resp., R_{Before} and R_{After}) denote the social surplus (resp., optimal revenue of the website) before and after the advent of ad-block software, respectively. We know that

$$R_{Before} = \begin{cases} \frac{Nr\bar{\rho}(V - V\alpha\rho - u_0)}{C_r}, & \text{if } u_0 \leq \hat{u}_0, \\ \\ \frac{Nr(V - u_0)}{C_b}, & \text{otherwise}, \end{cases} \tag{4.6}$$

and

$$R_{After} = \frac{Nr\bar{\rho}(V - u_0)}{C_r}. \tag{4.7}$$

The expressions for W_{Before} and W_{After} can be obtained using (4.4), (4.5), (4.6), and (4.7). The following result characterizes the movement in total surplus in the post-ad-block world.

Theorem 11. *After the advent of ad-block software, social surplus increases if and only if* $u_0 \leq \hat{u}_0$.

One direction of this result follows immediately from Theorems 8 and 10, since both the revenue of the website and the consumer surplus increase in the post-ad-block world when $u_0 \leq \hat{u}_0$. In the reverse direction (i.e., when $u_0 \geq \hat{u}_0$), while consumer surplus decreases in the post-ad-block world (Theorem 10), an attractive outside option for consumers compels the website to choose a low ad-intensity, thereby limiting the increase in its revenue. Consequently, the increase in the website's revenue is not enough to compensate for the loss of consumer surplus, leading to a net decrease in social surplus. Figure 4.5 summarizes the conclusions of Theorems 8, 10, and 11.

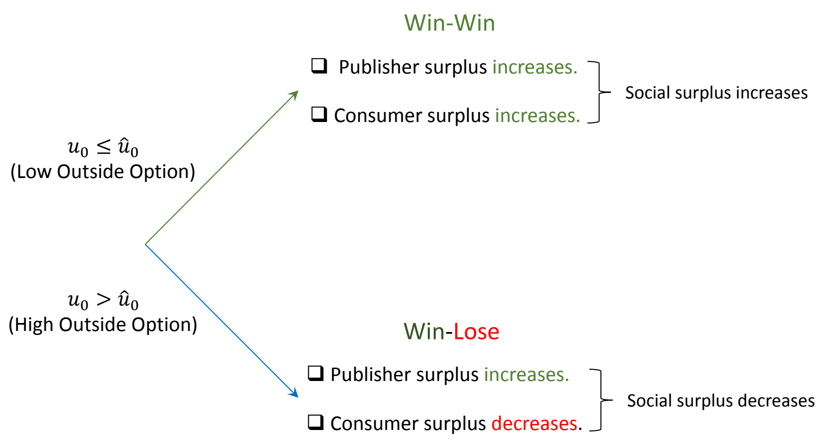

Figure 4.5. The advent of ad-blockers leads to a win-win for the websites that serve niche content ($u_0 \leq \hat{u}_0$) and their users.

In summary, relative to the pre-ad-block world, regular users are worse-off in the post-ad-block world while ad-block users (white-listers) are better-off (since they receive an ad-light experience). As long as the outside option is low enough, both the consumer surplus and the social surplus improve in the post-ad-block world (Theorem 10 and 11), and the website's revenue also improves (Theorem 8).

The expressions for R_{Before} and R_{After} in (4.6) and (4.7) help us establish the following result.

Theorem 12. *The rate of increase in the optimal revenue with an increase in the maximum value V of the website is higher in the post-ad-block world, i.e.,* $\frac{\partial R_{After}}{\partial V} \geq \frac{\partial R_{Before}}{\partial V}$*. Further, equality holds for $B = 0$ and $B = 1$.*

In both the pre- and post-ad-block worlds, the increase in the website's value from an increase in V enables it to advertise at a higher ad-intensity. However, because of the discriminatory power endowed by ad-blockers in the post-ad-block world, the website can customize the ad-intensity for each of the two user segments, i.e., regulars and ad-blockers. This is in contrast to the pre-ad-block world, where a common ad-intensity had to be chosen for the entire user population. This added advantage results in a relatively higher rate of increase in revenue with an increase in V in the post-ad-block world. Clearly, the ability to individually customize the ad intensities for the two segments can lead to a significant benefit only if both the segments are of sufficient size; i.e., the ad-blocking rate B is sufficiently away from 0 and 1.

A website can increase V by improving the quality of its content or service and/or by providing niche content. Thus, Theorem 12 implies that the website has a higher incentive to undertake such endeavors post the advent of ad-block software. Further, current ad-blocking rates are fractional for most websites, thus making the effort to increase V more potent. Finally, from (4.4) and (4.7), we also know that both consumer surplus and the website's revenue increase with V, leading to an increased social surplus. *In summary, content managers can view the provision of better/niche content or service as a curative response to ad-block software – one that is beneficial both to the website and its potential consumers and, therefore, generates more social value.*

Remark 3. *(White-Listing Cost)* Our analysis assumes that ad-block users do not incur any cost for white-listing the website. However, it is conceivable that users incur or perceive a positive *white-listing cost* – for example, users may feel annoyed when asked to white-list despite having an ad-blocker installed or non-tech-savvy users may find it difficult to properly white-list the website. In practice, websites try to reduce this cost by providing instructions on how to white-list and by explaining to users that ads are necessary to finance the creation of high-quality content. As long as the white-listing cost is sufficiently low, it is easy to show that the main results of our analysis – that the website's revenue increases after the advent of ad-blockers and that consumers also benefit under a low outside option – continue to hold. □

Our analysis thus far assumes that the website's gating decision is binary – either allow ad-free access to *all* ad-blocker users or require *all* these users to white-list to gain access. With the aim of further increasing the website's revenue, we now examine a generalization in which only a fraction of the ad-block users are gated.

4.5 Improving Revenue Further: Selective Gating

Let p denote the fraction of ad-blockers who are gated. The remaining $1 - p$ fraction of ad-blockers are not gated, i.e., are allowed to access the website without disabling the ad-block software. We refer to p as the *gating intensity* and to this strategy as *selective gating*. There are several ways in which selective gating can be implemented – e.g., tossing a coin (with success probability p) at each user visit to decide whether or not to gate, or randomly choosing a p fraction of users for gating. We will refer to the strategy in Section 4.2 of gating all or none of the potential users as *integral gating*, to contrast it with selective gating. Under selective gating, the sequence of events shown in Figure 4.2 of Section 4.1 will only change as follows: the indicator $I_G = 1$ is replaced by p and the indicator $I_G = 0$ is replaced by $(1 - p)$. Clearly, the equilibrium value of the website v will now be different from that in the case of integral gating (i.e., $I_G \in \{0, 1\}$). Thus, the decisions of the users change through the change in v. A regular user with an ad-sensitivity of $\tilde{c}_r$ becomes a user of the website if $v - a_r\tilde{c}_r \geq u_0$. An ad-blocker with an ad-sensitivity of $\tilde{c}_b$ becomes a user of this website under the following conditions: (a) if this user is selected for gating (with probability p) and $v - a_b\tilde{c}_b \geq u_0$, (b) if this user is not selected for gating (with probability $1-p$) and $v \geq u_0$; in this case, $\tilde{c}_b$ is inconsequential.

Our goal in this section is to identify the conditions under which the use of selective gating can lead to a substantial increase in the website's revenue, relative to integral gating. Recall from Theorem 4 that our analysis in Section 4.2 partitioned the website's decisions into three cases, based on the ad-viewing cost parameters C_b and C_r: (i) $C_b,\ C_r < \widehat{C}$, (ii) $C_b,\ C_r > \widehat{C}$, and (iii) $C_r \leq \widehat{C} \leq C_b$. In the first case, using an analysis similar to that in the proof of Theorem 4, we have $p^* = 1$, i.e., the optimal gating intensity under selective gating is naturally integral. Thus, in this case, selective gating does not affect the revenue of the website.

From the other two cases (namely, ii and iii), for brevity, we only discuss Case ii in this section, namely $C_b,\ C_r > \widehat{C}$. For this case, using an analysis similar to that in the proof of Theorem 4, we can show that, for any gating intensity $p \in [0,1]$, $a_b^* = a_r^* = a_{\min}$. Thus, we only need to obtain the optimal gating intensity p^*. Let $R(p, a_{\min}, a_{\min})$ represent the revenue of the website when gating intensity is p. Therefore, the website solves the following problem:

$$\max \quad R(p, a_{\min}, a_{\min}) \quad \text{s.t. } p \in [0,1]. \tag{P_{Select}}$$

Let p^* be the optimal gating intensity and let $v^* := v(p, a_{\min}, a_{\min})$. Let $n(p^*, a_{\min}, a_{\min})$ represent the equilibrium traffic received by the website. Similar to Sections 4.2.1 and 4.2.2, we first present some important properties that an optimal solution must satisfy and then use them to derive the optimal solution.

- **Property C1: (Rational Belief on Value)** The optimal gating intensity p^* is such that the equilibrium value v^* satisfies $0 \leq v^* - u_0 \leq a_{\min}C_r$.

- **Proof:** (i) If $v^* < u_0$, then the website does not receive any traffic and therefore does not earn any revenue. Therefore, we should have $v^* \geq u_0$. (ii) We know that $C_r \geq \widehat{C}$, where, from Section 4.2.1, $\widehat{C} = \frac{V-u_0}{a_{\min}}$. Since $V \geq v^*$, we have $C_r a_{\min} \geq v^* - u_0$. □

- **Property C2: (Equilibrium Traffic)** The equilibrium number of users is

$$n(p^*, a_{\min}, a_{\min}) = \frac{NB(v^* - u_0)p^*}{a_{\min} C_b} + (1-p^*)NB + \frac{N(1-B)(v^* - u_0)}{a_{\min} C_r}. \tag{4.8}$$

  - **Proof:** This result follows directly from Property C1 and our model of uniformly distributed ad-viewing costs for ad-blockers and regulars. □

The revenue function $R(p, a_{\min}, a_{\min})$ can now be computed using (4.8). It is easy to show that $R(p, a_{\min}, a_{\min})$ is a concave function of p. When $R'(0) \geq 0$ and $R'(1) \leq 0$, let $\hat{p}$ be such that $R'(\hat{p}) = 0$; clearly, such a value exists. We formally state the optimal gating intensity.

Theorem 13. *The optimal gating intensity is given by*

$$p^* = \begin{cases} 0, & \text{if } R'(0) < 0,\ R'(1) < 0, \\ 1, & \text{if } R'(0) > 0,\ R'(1) > 0, \\ \hat{p}, & \text{if } R'(0) \geq 0,\ R'(1) \leq 0. \end{cases}$$

We now proceed to examine the benefit of selective gating for the website.

4.5.1 Illustrating the Benefit of Selective Gating

To assess the benefit of selective gating, it is instructive to first understand the behavior of the optimal gating intensity p^*. We illustrate this using the following choice of the parameters: $C_r = 100,\ C_b = 200,\ u_0 = 50,\ a_{\min} = 1,\ N = 1{,}000{,}000,\ r = 1,$ and $V = 100$.

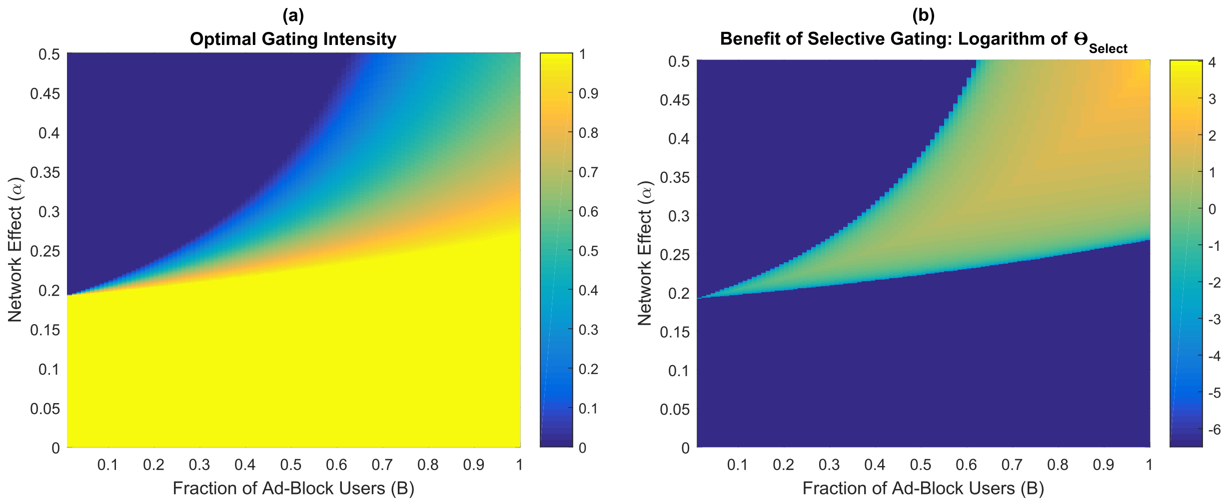

Figure 4.6. The benefit of selective gating is zero when the optimal gating intensity p^* is naturally integral, and is high for those combinations of α (the strength of the network effect) and B (the fraction of ad-block users) for which p^* is close to 0.5.

Figure 4.6(a) plots the optimal gating intensity (p^*) with respect to the network effect α and the ad-blocking rate B. Our main observations – which can be viewed as recommendations on the optimal gating intensity for the website based on the strength of its network effect and the ad-blocking rate of its user-population – are discussed below:

- From the lower half of the figure, we note the following:
  - When the strength of the network effect is low, it is optimal to always gate ad-blockers for all values of the ad-blocking rate, i.e., p^* is 1 for all values of B.
  - As the ad-blocking rate increases, the region in which p^* is 1 expands.

  The gating of ad-blockers affects a website in two ways: On the one hand, if an ad-blocker agrees to white-list the website, then revenue may be generated from this user who is otherwise a free-rider. On the other hand, if the ad-blocker does not white-list and leaves, then the resulting negative externality can adversely affect the user base of the website – the extent of this externality depends on the strength of the network effect for the website. Thus, when the magnitude of the network effect is low, the latter force is weak relative to the former and, as a result, it becomes optimal to always gate ad-blockers. Moreover, as the ad-blocking rate increases and ad-blockers constitute a larger fraction of the user-base, the conversion of ad-blockers to white-listers becomes even more attractive for the website, resulting in the expansion of the region in which p^* equals 1.

- The upper-left quadrant shows that a near-zero gating intensity is optimal when the ad-blocking rate is low but the network effect is strong. Both these features act in tandem towards a low gating intensity: while the strong network effect dissuades the website from gating ad-blockers, the low ad-blocking rate means that there is also no compelling need to react to ad-blockers.

- In the upper-right quadrant – where the network effect is strong and the ad-blocking rate is high – a fractional gating intensity is optimal. Here the tradeoff is healthy: while the strong network effect keeps the website from gating ad-blockers aggressively, the high ad-blocking rate also necessitates some action towards ad-blockers. This creates a fertile situation for selective gating to realize its potential, resulting in a fractional value of p^*. Further, as expected, p^* reduces as the network effect becomes stronger.

Another interesting observation from Figure 4.6(a) is that when the ad-blocking rate is low, the optimal gating intensity is largely integral, i.e., p^* is either 0 or 1. This perhaps explains why selective gating is not yet common in practice: Ad-blocking rates are still quite modest; in the U.S., for example, it is currently around 18%. As ad-blockers become prevalent, one can expect to see an increase in the use of the selective-gating strategy.

Benefit of Selective Gating Over Integral Gating: Having obtained the website's optimal revenue under both integral (Section 4.2) and selective gating, we can now examine the percentage increase in revenue from the latter. For our chosen values of the parameters, it can be verified that the optimal ad-intensities are $a_b^* = a_r^* = a_{\min}$ under both integral and selective gating. Let

$$\Theta_{Select} = \frac{R(p^*, a_{\min}, a_{\min}) - R(I_G^*, a_{\min}, a_{\min})}{R(I_G^*, a_{\min}, a_{\min})} \times 100.$$

Figure 4.6(b) plots the logarithm (base 10) of Θ_{Select} with respect to the network-effect parameter α (strength of the network effect) and the ad-blocking rate B (fraction of ad-block users). We now discuss our observations in this figure using additional information from Figure 4.6(a). Obviously, the benefit of selective gating is zero where the optimal gating intensity p^* is naturally integral. The benefit is highest in the region where p^* is around 0.5 (and therefore farthest from the integral extreme values). This is precisely the region in which both α and B are high. As is clear from our discussion above, concerns over the use of ad-blockers are particularly important for websites that face both a high network effect and a high fraction of ad-block users – these are the websites that can benefit most from selective gating.

Improvement in Revenue: In Theorem 8, we established that, under the optimal decisions derived in Section 4.2, the revenue of the website increases after the advent of ad-blockers. While that result assumed an integral gating decision, selective gating further increases the website's revenue. Recall from Section 4.2.2 that the website's revenue in the pre-ad-block world corresponds (in the post-ad-block world) to the website gating all potential users and using a common ad intensity $a_b = a_r = a$. Under our chosen values of the parameters, using Theorems 6 and the discussion earlier in this section, the website's revenue in the post-ad-blocker (resp., pre-ad-blocker) world is $R(p^*, a_{\min}, a_{\min})$ $\big($resp., $R(1, a_{\min}, a_{\min})\big)$. Thus, the percentage increase in the website's revenue after the advent of ad-blockers is

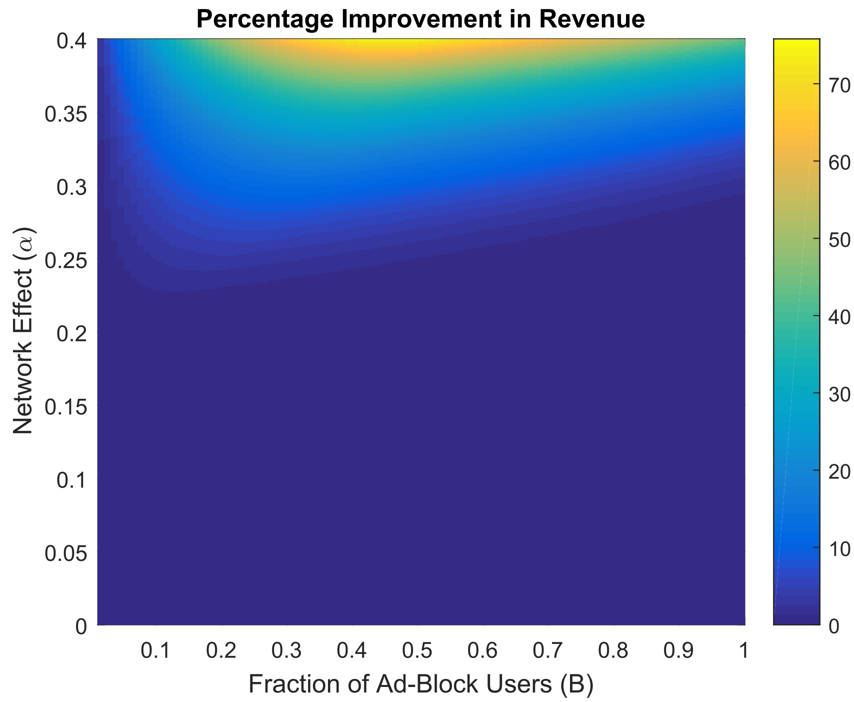

Figure 4.7. When the network effect is strong, the percentage increase in the website's revenue in the post-ad-blocking world peaks at a moderate ad-blocking rate B. When the network effect is weak, there is little or no increase in revenue.

$$\Upsilon_{Rev} = \frac{R(p^*, a_{\min}, a_{\min}) - R(1, a_{\min}, a_{\min})}{R(1, a_{\min}, a_{\min})} \times 100.$$

Figure 4.7 plots Υ_{Rev} with respect to ad-blocking rate B and network effect α. When the network effect is weak, the optimal gating intensity in the post-ad-block world is naturally close to 1 $\big($recall our earlier discussion for the lower half of Figure 4.6(a)$\big)$; therefore, there is little or no increase in revenue. When the network effect is strong, the percentage increase in revenue peaks for a moderate ad-blocking rate B (close to 0.5) – this is when the website can fully exploit the ability endowed by ad-blockers to discriminate between the two segments, namely regulars and ad-blockers.

4.6 Extensions: Robustness Check under Additional Features

We now analyze the following additional features as extensions to our base model in Section 4.1:

- Heterogeneous profitability of the two user groups – namely, ad-block users and regular users – to the website,
- The offering of a paid, ad-free subscription option by the website (in addition to the option of viewing ad-supported free content by white-listing the website),
- Endogenous adoption of ad-blockers by users in response to the website's strategy (i.e., its chosen ad-intensities for white-listers and regular users), and
- Negative externality (congestion cost) imposed on the website by an increase in traffic.

We add each of these features, one at a time, to our base model and establish the validity of our main conclusions: (i) the revenue of the website increases in the post-ad-block world relative to that in the pre-ad-block world (Theorem 8), (ii) ad-block users receive an ad-light experience (Theorem 9), and (iii) the social surplus can either increase or decrease in the post-ad-block world relative to that in the pre-ad-block world (Theorem 11).

Recall that, in our base model, the domain for the ad-intensities a_r and a_b (decisions of the website) for, respectively, regular users and white-listers is $\{0\} \cup [a_{\min}, \infty]$. That is, a_r (resp., a_b) is either 0 or at least $a_{\min}$. For simplicity of exposition, we present the analysis of the above extensions under the assumption $a_{\min} = 0$.

In the following subsections, we define the setting for each of the four extensions; the corresponding technical analysis is provided in Appendix.

4.6.1 Heterogeneous Profitability

Our basic model in Section 4.1 assumes that both ad-block users and regular users are equally profitable. That is, the website earns the same revenue r per unit ad-intensity from both ad-block users and regular users. We now relax this assumption and let these two user groups be heterogeneous in terms of their profitability to the website. Specifically, we assume that the website earns a revenue of r_b per unit ad-intensity from ad-block users, and a revenue of r_g per unit ad-intensity from regular users. We make no assumption on the relative comparison of these two values (i.e., our analysis holds regardless of whether $r_b \leq r_g$ or $r_b \geq r_g$). The analysis of this extension is in the Appendix.

4.6.2 Subscription Option

Our basic model of Section 4.1 considers a website that operates only on advertising revenue. While this is indeed the case for most websites on the internet, there are websites that also offer a paid subscription option to users for viewing ad-free content. To ad-block users, such websites typically give an additional option of subscription after gating them. In this extension, we model a website that offers its users both ad-supported free content and ad-free paid content (subscription).

We assume that the website charges a subscription fee f per visit (a decision for the website) for its ad-free content (if the subscription fee is charged periodically, say per year, then it can be appropriately amortized to a per-visit fee). Figure 4.8 depicts the sequence of events when the website offers the subscription option. If the website decides to gate an ad-block user, then that user has the following three options: (i) white-list the website, (ii) subscribe to the website, or (iii) leave the website and consume the outside option. The website decides the optimal subscription fee f along with the optimal ad-intensities a_b and a_r. In the pre-ad-block world, the website decides the optimal subscription fee f and the optimal ad-intensity a for the entire user population. The analysis of this extension is in the Appendix.

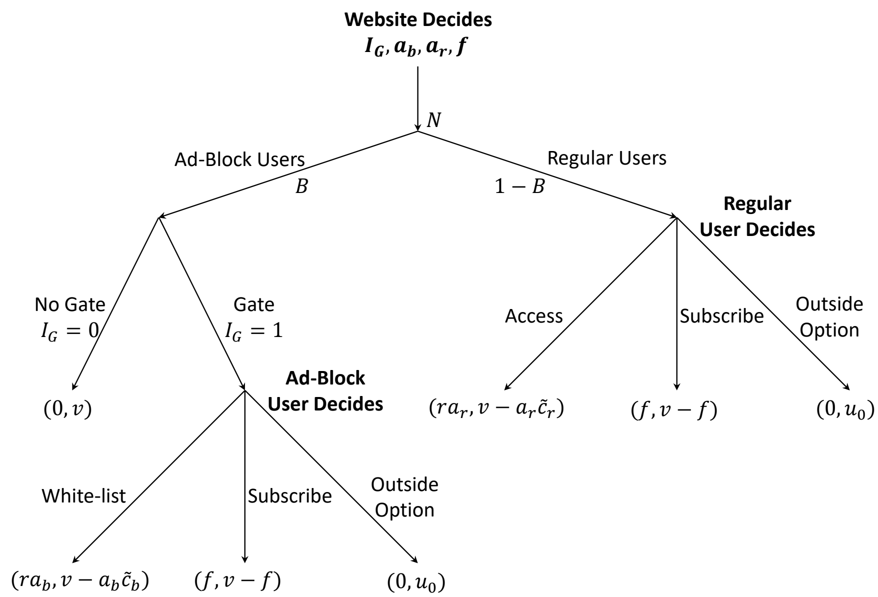

Figure 4.8. The sequence of events of the game when the website offers both the white-listing and subscription options: The website first decides I_G, a_b, a_r, and f. The ad-blockers then decide whether to white-list (for ad-supported content) or subscribe (for ad-free content) or consume the outside option. Similarly, the regulars decide whether to access the website with ads or subscribe or consume the outside option. The quantities in brackets represent the payoffs received by the website and the user, respectively.

4.6.3 Endogenous Adoption of Ad-Blockers

In our basic model of Section 4.1, we consider two exogenously given groups of users (ad block users and regular users). Instead, in this extension, we let these two groups form endogenously in response to the website's strategy (i.e., its chosen ad-intensities for white-listers and regular users). In practice, the adoption of ad-blockers by users is typically influenced by the ad-intensities of a multitude of websites. Thus, the setting where a single website can control the adoption of ad-blockers by changing its ad-intensities is realistic only for influential websites that serve a large number of users (e.g., Youtube.com).

We consider such a website in this extension. For our analysis in this section, it is convenient to view v as the (infinite-horizon, discounted) lifetime value of the website to a user. We assume a potential user population with heterogenous ad-sensitivities: For a randomly-picked potential user who is shown an ad-intensity of 1, the ad-viewing cost is a random variable denoted by $\tilde{c}_a \sim U[0,\ C_a]$; a similar lifetime interpretation is convenient for these costs as well as the maximum-possible value V from the website. We assume that a user incurs a cost $S \geq 0$ in installing an ad-blocker. The analysis of this extension is in the Appendix.

4.6.4 Negative Externality (Congestion Cost) Due to Increase in Traffic

Our basic model of Section 4.1 does not consider any negative externality imposed on the website due to an increase in traffic. In this extension, the website incurs a congestion cost due to increased traffic. Specifically, for a constant $k_c \geq 0$, we assume that the website incurs a congestion cost of $k_c n^2$, where n is the equilibrium traffic of the website. It is convenient to interpret this as a "service cost"; that is, the cost incurred by the website in increasing server capacity (due to increased traffic) to maintain a fixed service level. The analysis of this extension is in the Appendix.

Next, we discuss the recommendations offered by our analysis for publishers, web-browsers, and policy makers.

4.7 Recommendations for Publishers, Browsers, and Policy Makers

Publishers: While publishers always benefit in the post-ad-block world (Section 4.3), the specific strategy – namely, the gating decision and the ad-intensities for regulars and ad-blockers – to realize this benefit differs across publishers. Two key characteristics that

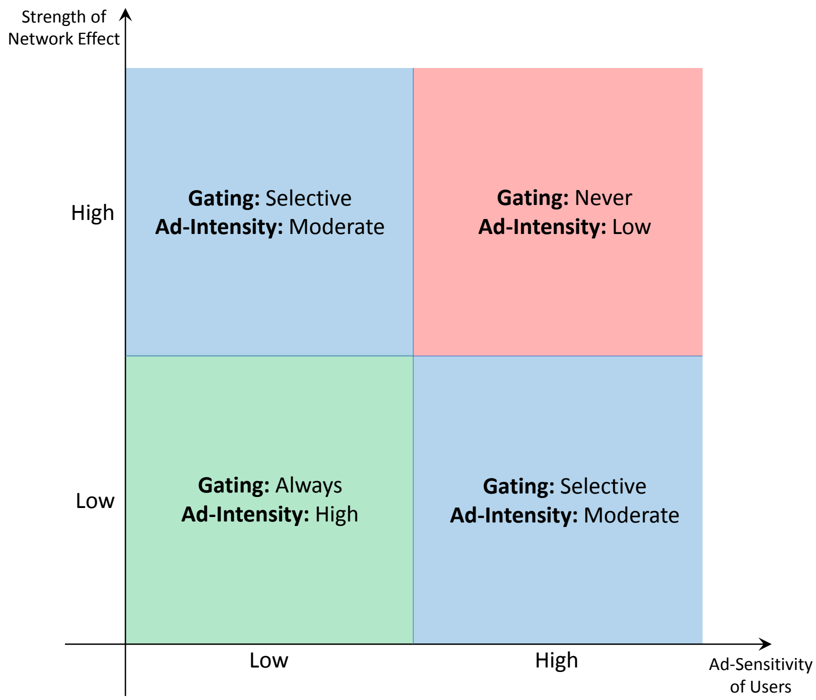

Figure 4.9. The recommended gating decision and advertisement intensity based on the strength of a website's network effect and the ad-sensitivity of its potential users.

shape a publisher's strategy are the strength of the network effect it faces and the ad-sensitivity of its potential users. Figure 4.9 summarizes the associated recommendations offered by our analysis: If the network effect is strong and the user population is ad-sensitive, then no gating and a low ad-intensity is recommended. On the other extreme, if the network effect is low and users have a high ad-tolerance, then always gating and a high ad-intensity is suggested. Otherwise, the website should selectively gate users and advertise at a moderate intensity. In all these cases, when a website gates, it is recommended that white-listers be offered an ad-light experience.

Browsers: Recently, Google announced that it plans to provide an inbuilt ad-blocker for its Chrome web browser (Wall Street Journal, 2017). There already exist other web browsers that come with an inbuilt ad-blocker, e.g., Adblock Browser for mobiles, and Opera for both desktops and mobiles. We now discuss the potential impact of the decision of web browsers to provide an inbuilt ad-blocker on the welfare of publishers and consumers. Specifically, let us imagine the following setting: The web browsers of *all potential users* of the website have an inbuilt ad-blocker with a white-listing feature (both Opera and Adblock Browser offer this feature). Thus, the only way that the website can show ads to its users is by asking them to white-list the website on their browsers.

As our analysis in Sections 4.2 and 4.3 showed, the conscious decision of users to install an ad-blocker divulges information about their ad-viewing costs, which can then be exploited by the website to tailor their ad-intensities. When the outside option offers low utility, this separation not only benefits publishers but also improves consumer surplus. On the flip side, this discrimination is not powerful enough when the outside option offers a high-enough utility, in the sense that users with low ad-viewing costs can be worse-off after the advent of ad-blockers. However, if *all* potential users have an inbuilt ad-blocker, the publisher cannot distinguish them on the basis of their ad-viewing costs and therefore loses the ability to customize their ad-intensities. Thus, we are back to the situation *before* the advent of ad blockers. Let $P_{Inbuilt}$ refer to the website's problem for a user population with in-built ad-blockers. The analysis of this problem is identical to that in Section 4.2.2 of problem P_{Before} – the website will gate all users and offer them a single ad-intensity upon white-listing. Thus, the comparison between $P_{Inbuilt}$ and P_{After} is identical to the comparison between P_{After} and P_{Before} with the trends *reversed*. That is, with in-built ad-blockers (i) the website's revenue decreases and (ii) both the consumer surplus and social surplus increase if $u_0 \geq \hat{u}_0$, and decrease otherwise.

It would be safe to assume that the primary goal of a web browser is to provide a compelling browsing experience to its users. Thus, the decision of web browsers to provide an inbuilt ad-blocker is perhaps aimed at improving consumer surplus. However, our analysis suggests that inbuilt ad-blockers increase consumer welfare of only those websites for which the outside option is high ($u_0 \geq \hat{u}_0$). Publishers that offer unique content – e.g., Facebook, Twitter, and Youtube – arguably have low-utility outside options; their users can be worse-off with in-built ad-blockers. Overall, we conclude that the well-intentioned intervention to provide inbuilt ad-blockers can have undesirable consequences and should therefore be maneuvered carefully. Alternately, it would be best to let the ecosystem evolve organically, i.e., let users consciously make a choice to install ad-blockers.

Policy Makers: Publishers and firms that develop ad-block software continue to engage in protracted legal battles. With ad-revenue being their only source of revenue, publishers have argued for the imposition of a legal ban on ad-block software – a charge that ad-blocking firms have thus far successfully fought against (Meyer, 2016; Rodriguez, 2017). Also, privacy activists have objected to publishers detecting ad-blockers in users' web browsers without their permission (Heilpern, 2016). Our analysis derives the conditions under which both publishers and users can benefit from ad-block software and also derives the associated operational decisions for the publisher. Perhaps these results can help the stakeholders better understand the implications of ad-block software and, consequently, de-escalate the conflict between the two parties.

4.8 Concluding Remarks

A typical website consists of multiple webpages. An underlying assumption of our analysis in this paper was that the website's decisions, namely those of gating and choosing the ad-intensities for regulars and ad-blockers, are the same for each page of the website. It is not uncommon, however, for some pages of a website to offer unique content (for instance, an opinion piece on nytimes.com) while other pages provide generic content that is easily available elsewhere. Naturally, it is easier for the website to convince potential users to whitelist pages that offer unique content. Thus, one can potentially investigate a richer set of decisions for the website; e.g., an aggressive gating policy for pages with unique content and a relatively relaxed one for those with generic content. This setting can be challenging to analyze, for example due to the subtle interaction between the equilibrium traffic on these two types of pages. Another potential enhancement to our setting would be to incorporate the presence of a competing website.

APPENDIX

PROOF OF THEOREMS

Proof of Theorem 1: Consider any policy π which is feasible to $\mathcal{P}^E_{static}(\vec{\beta})$. Define the vector $(\tilde{x}, \tilde{y})$ as follows:

$$\tilde{x}_{s,i,l} = \frac{\sum_{t\in\mathcal{T}_{s,i}} \mathbb{E}_{\vec{U}_{t-1}}\left[x^{\pi}_{t,l}\left(\vec{j}^{\pi}_t(\vec{U}_{t-1})\right)\right]}{\Delta}, \tag{A.1}$$

$$\tilde{y}_{s,i,l,c} = \frac{\sum_{t\in\mathcal{T}_{s,i}} \mathbb{E}_{\vec{U}_{t-1}}\left[y^{\pi}_{t,l,c}\left(\vec{j}^{\pi}_t(\vec{U}_{t-1})\right)\right]}{\Delta}. \tag{A.2}$$

Our proof involves demonstrating the following claims:

(a) $(\tilde{x}, \tilde{y})$ is feasible to $\hat{\mathcal{P}}(\vec{\beta})$.

(b) $\pi(\vec{\beta})$ is feasible to $\mathcal{P}^E_{static}(\vec{\beta})$.

(c) $f^{\pi} \geq \hat{f}(\tilde{x})$.

(d) $\hat{f}(\tilde{x}) \geq \hat{f}\left(\hat{x}^*(\vec{\beta})\right)$.

(e) $\hat{f}\left(\hat{x}^*(\vec{\beta})\right) = f^{\pi(\vec{\beta})}$.

Combining (a) – (e) we obtain the desired result that $f^{\pi(\vec{\beta})} \leq f^{\pi}$ for any policy π which is feasible for $\mathcal{P}^E_{static}(\vec{\beta})$. Thus, $\pi(\vec{\beta})$ is optimal for $\mathcal{P}^E_{static}(\vec{\beta})$. It only remains to show (a) – (e).

Proof of (a): The feasibility of $(\tilde{x}, \tilde{y})$ for $\hat{\mathcal{P}}(\vec{\beta})$ follows from its definition and the feasibility of policy π for problem $\mathcal{P}^E_{static}(\vec{\beta})$.

Proof of (b): The feasibility of $\pi(\vec{\beta})$ for problem $\mathcal{P}^E_{static}(\vec{\beta})$ follows from its definition and the feasibility of $\left(\hat{x}^*(\vec{\beta}),\ \hat{y}^*(\vec{\beta})\right)$ for problem $\hat{\mathcal{P}}(\vec{\beta})$.

Proof of (c): The following two inequalities follow from Jensen's inequality:

$$\mathbb{E}_{\vec{U}_{t-1}}\left[f_{i,l}\left(x^{\pi}_{t,l}\left(\vec{j}^{\pi}_t(\vec{U}_{t-1})\right)\right)\right] \geq f_{i,l}\left(\mathbb{E}_{\vec{U}_{t-1}}\left[x^{\pi}_{t,l}\left(\vec{j}^{\pi}_t(\vec{U}_{t-1})\right)\right]\right),\ \forall\, t \in \mathcal{T}_{s,i},\ s \in \mathcal{S},\ i \in \mathcal{I},\ l \in \mathcal{L},$$

and

$$\frac{\sum_{t\in\mathcal{T}_{s,i}} \mathbb{E}_{\vec{U}_{t-1}}\Big[f_{i,l}\Big(x^{\pi}_{t,l}\big(\vec{j}^{\pi}_t(\vec{U}_{t-1})\big)\Big)\Big]}{\Delta} \geq f_{i,l}\left(\frac{\sum_{t\in\mathcal{T}_{s,i}} \mathbb{E}_{\vec{U}_{t-1}}\big[x^{\pi}_{t,l}\big(\vec{j}^{\pi}_t(\vec{U}_{t-1})\big)\big]}{\Delta}\right)$$

$$= f_{i,l}(\tilde{x}_{i,l}),\ \forall\, s\in\mathcal{S},\ i\in\mathcal{I},\ l\in\mathcal{L}.$$

Multiplying both sides by $\Delta q_{i,l}$ and summing over $l\in\mathcal{L}$, we obtain

$$\sum_{t\in\mathcal{T}_{s,i}}\sum_{l\in\mathcal{L}} q_{i,l}\mathbb{E}_{\vec{U}_{t-1}}\Big[f_{i,l}\Big(x^{\pi}_{t,l}\big(\vec{j}^{\pi}_t(\vec{U}_{t-1})\big)\Big)\Big] \geq \Delta\sum_{l\in\mathcal{L}} q_{i,l} f_{i,l}(\tilde{x}_{i,l}).$$

Summing over $s\in\mathcal{S},\ i\in\mathcal{I}$, we have

$$f^{\pi} \geq \hat{f}(\tilde{x}).$$

Proof of (d): This claim follows immediately from claim (a) and the optimality of $\hat{x}^*(\vec{\beta})$ for problem $\hat{\mathcal{P}}(\vec{\beta})$.

Proof of (e): Since $\pi(\vec{\beta})$ is a state independent policy, directly substituting its definition in the expression f^{π} gives $f^{\pi(\vec{\beta})} = \hat{f}\big(\hat{x}^*(\vec{\beta})\big)$. ∎

Proof of Theorem 2: We want to find a bound for the ratio

$$\frac{f^{\pi(\vec{\beta}^*)}}{\text{Opt}(\vec{M},\alpha)}\ .$$

Recall that problem $\mathcal{P}^E_{static}(\alpha\vec{M})$ is a relaxation of problem $\mathcal{P}_{static}(\vec{M},\alpha)$ and that $\hat{\mathcal{P}}(\alpha\vec{M})$ is equivalent to $\mathcal{P}^E_{static}(\alpha\vec{M})$. Let $\pi(\alpha\vec{M}) = \big(\hat{x}^*(\alpha\vec{M}),\ \hat{y}^*(\alpha\vec{M})\big)$ be the solution of problem $\hat{\mathcal{P}}(\alpha\vec{M})$, where

$$\hat{x}^*(\alpha\vec{M}) = \big(\hat{x}^*_{s,i,l}(\alpha\vec{M}):\ s\in\mathcal{S},\ i\in\mathcal{I},\ l\in\mathcal{L}\big)$$

and

$$\hat{y}^*(\alpha\vec{M}) = \big(\hat{y}^*_{s,i,l,c}(\alpha\vec{M}):\ s\in\mathcal{S},\ i\in\mathcal{I},\ l\in\mathcal{L},\ c\in\mathcal{C}\big).$$

Therefore,

$$\text{Opt}(\vec{M},\alpha) \geq \hat{f}\big(\hat{x}^*(\alpha\vec{M})\big) = \sum_{s\in\mathcal{S}}\sum_{i\in\mathcal{I}}\sum_{l\in\mathcal{L}} \Delta\ q_{i,l}\ f_{i,l}\big(\hat{x}^*_{s,i,l}(\alpha\vec{M})\big). \tag{A.3}$$

We can obtain a feasible solution for problem $\hat{\mathcal{P}}(\vec{M} + z_\alpha \vec{M}_r)$, where $\vec{M}_r \in \mathbb{R}_+^{|\mathcal{C}|}$ is a vector with elements $\sqrt{M_c}, c \in \mathcal{C}$, by simply multiplying the solution $\big(\hat{x}^*(\alpha\vec{M}),\ \hat{y}^*(\alpha\vec{M})\big)$ of problem $\hat{\mathcal{P}}(\alpha\vec{M})$ by the factor

$$\gamma = \max_{c\in\mathcal{C}} \left[\frac{M_c + z_\alpha\sqrt{M_c}}{\alpha M_c}\right] = \max_{c\in\mathcal{C}} \left[\frac{1}{\alpha}\left(1 + \frac{z_\alpha}{\sqrt{M_c}}\right)\right].$$

Thus, $\pi^\gamma(\alpha\vec{M}) := \big(\gamma\hat{x}^*(\alpha\vec{M}),\ \gamma\hat{y}^*(\alpha\vec{M})\big)$ is a feasible solution for problem $\hat{\mathcal{P}}(\vec{M} + z_\alpha\vec{M}_r)$.

We know that the cost associated with policy $\pi^\gamma(\alpha\vec{M})$ is $f^{\pi^\gamma(\alpha\vec{M})}$. Since $\pi^\gamma(\alpha\vec{M})$ is only a feasible policy for problem $\hat{\mathcal{P}}(\vec{M} + z_\alpha\vec{M}_r)$, the cost corresponding to $\pi^\gamma(\alpha\vec{M})$, i.e., $f^{\pi^\gamma(\alpha\vec{M})}$, is higher than the optimal cost for problem $\hat{\mathcal{P}}(\vec{M} + z_\alpha\vec{M}_r)$, i.e., $f^{\pi(\vec{\beta}^*)}$. Thus, $f^{\pi(\vec{\beta}^*)} \leq f^{\pi^\gamma(\alpha\vec{M})}$. We also know that $f^{\pi^\gamma(\alpha\vec{M})} = \sum_{s\in\mathcal{S}}\sum_{i\in\mathcal{I}}\sum_{l\in\mathcal{L}} \Delta\ q_{i,l}\ f_{i,l}\big(\gamma\hat{x}^*_{s,i,l}(\alpha\vec{M})\big)$. Therefore,

$$f^{\pi(\vec{\beta}^*)} \leq \sum_{s\in\mathcal{S}}\sum_{i\in\mathcal{I}}\sum_{l\in\mathcal{L}} \Delta\ q_{i,l}\ f_{i,l}\big(\gamma\hat{x}^*_{s,i,l}(\alpha\vec{M})\big). \tag{A.4}$$

From (A.3) and (A.4) we can write,

$$\frac{f^{\pi(\vec{\beta}^*)}}{\text{Opt}(\vec{M},\alpha)} \leq \frac{\sum_{s\in\mathcal{S}}\sum_{i\in\mathcal{I}}\sum_{l\in\mathcal{L}} \Delta\ q_{i,l}\ f_{i,l}\big(\gamma\hat{x}^*_{s,i,l}(\alpha\vec{M})\big)}{\sum_{s\in\mathcal{S}}\sum_{i\in\mathcal{I}}\sum_{l\in\mathcal{L}} \Delta\ q_{i,l}\ f_{i,l}\big(\hat{x}^*_{s,i,l}(\alpha\vec{M})\big)}. \tag{A.5}$$

Now, consider the ratio,

$$\frac{f_{i,l}\big(\gamma\hat{x}^*_{s,i,l}(\alpha\vec{M})\big)}{f_{i,l}\big(\hat{x}^*_{s,i,l}(\alpha\vec{M})\big)}.$$

Since $\gamma \geq 1$, $\gamma\hat{x}^*_{s,i,l}(\alpha\vec{M}) \geq \hat{x}^*_{s,i,l}(\alpha\vec{M})$. Further, we know that $f_{i,l}(\cdot)$ is convex. Thus, we can write,

$$\begin{aligned}
\frac{f_{i,l}\big(\gamma\hat{x}^*_{s,i,l}(\alpha\vec{M})\big)}{f_{i,l}\big(\hat{x}^*_{s,i,l}(\alpha\vec{M})\big)} &\leq \frac{f_{i,l}\big(\gamma\hat{x}^*_{s,i,l}(\alpha\vec{M})\big)}{f_{i,l}\big(\gamma\hat{x}^*_{s,i,l}(\alpha\vec{M})\big) - \big(\gamma\hat{x}^*_{s,i,l}(\alpha\vec{M}) - \hat{x}^*_{s,i,l}(\alpha\vec{M})\big)f'_{i,l}\big(\gamma\hat{x}^*_{s,i,l}(\alpha\vec{M})\big)}, \\
&= \frac{f_{i,l}\big(\gamma\hat{x}^*_{s,i,l}(\alpha\vec{M})\big)}{f_{i,l}\big(\gamma\hat{x}^*_{s,i,l}(\alpha\vec{M})\big) - (\gamma-1)\hat{x}^*_{s,i,l}(\alpha\vec{M})f'_{i,l}\big(\gamma\hat{x}^*_{s,i,l}(\alpha\vec{M})\big)}, \\
&= \frac{1}{1 - \frac{(\gamma-1)}{\gamma}\frac{\gamma\hat{x}^*_{s,i,l}(\alpha\vec{M})f'_{i,l}\left(\gamma\hat{x}^*_{s,i,l}(\alpha\vec{M})\right)}{f_{i,l}\left(\gamma\hat{x}^*_{s,i,l}(\alpha\vec{M})\right)}}.
\end{aligned}$$

Now, the above inequality can be rewritten as

$$\frac{f_{i,l}\big(\gamma\hat{x}^*_{s,i,l}(\alpha\vec{M})\big)}{f_{i,l}\big(\hat{x}^*_{s,i,l}(\alpha\vec{M})\big)} \leq \frac{1}{1-\frac{(\gamma-1)}{\gamma}\psi_{i,l}\big(\gamma\hat{x}^*_{s,i,l}(\alpha\vec{M})\big)}, \\ \leq \frac{1}{1-\frac{(\gamma-1)}{\gamma}\overline{\psi}}. \tag{A.6}$$

Using (A.5) and (A.6), we obtain the desired result that

$$\frac{f^{\pi(\vec{\beta}^*)}}{\text{Opt}(\vec{M},\alpha)} \leq \frac{1}{1-\frac{(\gamma-1)}{\gamma}\overline{\psi}}.$$

This completes the proof of Theorem 2. ■

Proof of Proposition 1: The simple intuition behind the proof is as follows: Consider the set of active campaigns at location l at any point in time. As time progresses, some campaigns end, resulting in fewer campaigns being active in the later time blocks. Thus, a higher number of active campaigns in the earlier time blocks leads to higher target win-probabilities than those in the later time blocks.

Let $\hat{x}(\vec{\beta}) = \{\hat{x}_{s,i,l}(\vec{\beta}) :\ s \in \mathcal{S},\ i \in \mathcal{I},\ l \in \mathcal{L}\}$ and $\hat{y}(\vec{\beta}) = \{\hat{y}_{s,i,l,c}(\vec{\beta}) :\ s \in \mathcal{S},\ i \in \mathcal{I},\ l \in \mathcal{L},\ c \in \mathcal{C}\}$ be an optimal solution of problem $\hat{\mathcal{P}}(\vec{\beta})$. Consider an arbitrary time block $i \in \mathcal{I}$ and location $l \in \mathcal{L}$, and assume, for a contradiction, that there exist two consecutive time periods d and $d+1$ with $\hat{x}_{d,i,l}(\vec{\beta}) < \hat{x}_{d+1,i,l}(\vec{\beta})$. That is, in time block i at location l, the target win-probability in the d-th period is strictly less than the target win-probability in the $(d+1)$-th period. We will construct another feasible solution $\big(\hat{x}'(\vec{\beta}),\ \hat{y}'(\vec{\beta})\big)$ such that (i) $\hat{x}'_{d,i,l}(\vec{\beta}) = \hat{x}'_{d+1,i,l}(\vec{\beta})$, (ii) $\hat{x}'_{s,i,l}(\vec{\beta}) = \hat{x}_{s,i,l}(\vec{\beta}),\ s \in \mathcal{S} \setminus \{d, d+1\}$, (iii) $\hat{x}'_{s',i',l'}(\vec{\beta}) = \hat{x}_{s',i',l'}(\vec{\beta}),\ (s',i',l') \in \mathcal{S}\times\mathcal{I}\times\mathcal{L} \setminus \big\{(d,i,l),\ (d+1,i,l)\big\}$, and (iv) the objective function value of the solution $\big(\hat{x}'(\vec{\beta}), \hat{y}'(\vec{\beta})\big)$ is strictly less than that of the solution $\big(\hat{x}(\vec{\beta}), \hat{y}(\vec{\beta})\big)$.

Let

$$\bar{x}_{d,d+1,i,l}(\vec{\beta}) = \frac{\hat{x}_{d,i,l}(\vec{\beta}) + \hat{x}_{d+1,i,l}(\vec{\beta})}{2}\ .$$

For location l and time block i, define

$$\hat{x}'_{d,i,l}(\vec{\beta}) = \hat{x}'_{d+1,i,l}(\vec{\beta}) = \bar{x}_{d,d+1,i,l}(\vec{\beta}), \tag{A.7}$$
$$\hat{x}'_{s,i,l}(\vec{\beta}) = \hat{x}_{s,i,l}(\vec{\beta}),\ \forall\, s \in \mathcal{S} \setminus \{d,\ d+1\}.$$

For all other time periods, time blocks, and locations define

$$\hat{x}'_{s',i',l'}(\vec{\beta}) = \hat{x}_{s',i',l'}(\vec{\beta}),\ (s',i',l') \in \mathcal{S} \times \mathcal{I} \times \mathcal{L} \setminus \big\{(d,i,l),\ (d+1,i,l)\big\}.$$

We now define the allocation probabilities $\hat{y}'(\vec{\beta})$ by adjusting the original allocation probabilities $\hat{y}(\vec{\beta})$. We do this by advancing the procurment for some campaigns from time period $d+1$ to time period d. Let

$$c^* = \min\left\{\tilde{c} : \sum_{c=1}^{\tilde{c}} \hat{y}_{d+1,i,l,c}(\vec{\beta}) \geq \hat{x}_{d+1,i,l} - \bar{x}_{d,d+1,i,l}(\vec{\beta})\right\}.$$

For location l, define

$$\begin{aligned}
\hat{y}'_{d,i,l,c}(\vec{\beta}) &= \hat{y}_{d,i,l,c}(\vec{\beta}) + \hat{y}_{d+1,i,l,c}(\vec{\beta}), \qquad \forall\, c \in \{1,2,\ldots,c^*-1\},\\
\hat{y}'_{d+1,i,l,c}(\vec{\beta}) &= 0, \qquad \forall\, c \in \{1,2,\ldots,c^*-1\},\\
\hat{y}'_{d,i,l,c^*}(\vec{\beta}) &= \hat{y}_{d,i,l,c^*}(\vec{\beta}) + \left[\big(\hat{x}_{d+1,i,l}(\vec{\beta}) - \bar{x}_{d,d+1,i,l}(\vec{\beta})\big) - \sum_{c=1}^{c^*-1} \hat{y}_{d+1,i,l,c}(\vec{\beta})\right],\\
\hat{y}'_{d+1,i,l,c^*}(\vec{\beta}) &= \hat{y}_{d+1,i,l,c^*}(\vec{\beta}) - \left[\big(\hat{x}_{d+1,i,l}(\vec{\beta}) - \bar{x}_{d,d+1,i,l}(\vec{\beta})\big) - \sum_{c=1}^{c^*-1} \hat{y}_{d+1,i,l,c}(\vec{\beta})\right],\\
\hat{y}'_{d,i,l,c}(\vec{\beta}) &= \hat{y}_{d,i,l,c}(\vec{\beta}) \qquad \forall\, c \in \{c^*+1, c^*+2, \ldots, |\mathcal{C}|\},\\
\hat{y}'_{d+1,i,l,c}(\vec{\beta}) &= \hat{y}_{d+1,i,l,c}(\vec{\beta}) \qquad \forall\, c \in \{c^*+1, c^*+2, \ldots, |\mathcal{C}|\}.
\end{aligned}$$

For all other time periods, time blocks, and locations define

$$\hat{y}'_{s',i',l',c}(\vec{\beta}) = \hat{y}_{s',i',l',c}(\vec{\beta}),\ (s',i',l') \in \mathcal{S} \times \mathcal{I} \times \mathcal{L} \setminus \big\{(d,i,l),\ (d+1,i,l)\big\},\ c \in \mathcal{C}.$$

It is easy to verify that $\big(\hat{x}'(\vec{\beta}),\ \hat{y}'(\vec{\beta})\big)$ is a feasible solution of problem $\hat{\mathcal{P}}(\vec{\beta})$. The fact that the cost of $\big(\hat{x}'(\vec{\beta}),\ \hat{y}'(\vec{\beta})\big)$ is strictly less than that of $\big(\hat{x}(\vec{\beta}),\ \hat{y}(\vec{\beta})\big)$ is a direct consequence of the strict convexity of $f_{i,l}(\cdot)$. This completes the proof. ■

Proof of Proposition 2: Define $\hat{x}^*(\vec{\beta}) = \{\hat{x}^*_{s,i,l}(\vec{\beta}) : \ s \in \mathcal{S},\ i \in \mathcal{I},\ l \in \mathcal{L}\}$ and $\hat{y}^*(\vec{\beta}) = \{\hat{y}^*_{s,i,l,c}(\vec{\beta}) : \ s \in \mathcal{S},\ i \in \mathcal{I},\ l \in \mathcal{L},\ c \in \mathcal{C}\}$. Let $\big(\hat{x}^*(\vec{\beta}),\ \hat{y}^*(\vec{\beta})\big)$ be an optimal solution of problem $\hat{\mathcal{P}}(\vec{\beta})$. We want to prove that the optimal win-probability in location l during time block i of time period s decreases with an increase in the impression-arrival probability $q_{i,l}$. That is, we want to show that $\hat{x}^*_{s,i,l}(\vec{\beta})$ decreases with $q_{i,l}$. Without loss of generality, we prove this for $\hat{x}^*_{1,1,1}(\vec{\beta})$, i.e., we prove that $\hat{x}^*_{1,1,1}(\vec{\beta})$ decreases with an increase in $q_{1,1}$. Define $\mathcal{B}_{-1} = \mathcal{S} \times \mathcal{I} \times \mathcal{L} \setminus (1,1,1)$. Let $\hat{x}_{-1}(\vec{\beta}) = \big\{\hat{x}_{s,i,l}(\vec{\beta}) : \ (s,i,l) \in \mathcal{B}_{-1}\big\}$. We now define a new problem in which for a given value of win-probability in location $l = 1$ for time block $i = 1$ and time period $s = 1$, we find all other win-probabilities. For a given value of $\hat{x}_{1,1,1}$, we define $z = \Delta q_{1,1} \hat{x}_{1,1,1}$.

Thus, we solve the following optimization problem to obtain the value of objective function at the optimal values of the other win-probabilities, given $\hat{x}_{1,1,1}$:

$$
\hat{\mathcal{P}}_{-1}(\vec{\beta}) : \left\{
\begin{aligned}
& \mathcal{F}(z) = \min_{\hat{x}_{-1}, \hat{y}} \sum_{(s,i,l) \in \mathcal{B}_{-1}} \Delta\ q_{i,l}\ f_{i,l}(\hat{x}_{s,i,l}) \\
& \text{subject to:} \\
& \sum_{c \in \mathcal{C}} \hat{y}_{s,i,l,c} \ = \ \hat{x}_{s,i,l},\ \forall\ (s,i,l) \in \mathcal{B}_{-1}, && \text{(A.8)} \\
& \Delta q_{1,1} \sum_{c \in \mathcal{C}} \hat{y}_{1,1,1,c} \ = \ z, \\
& \sum_{(s,i,l) \in \mathcal{B}_{-1}} \Delta\ q_{i,l} \hat{y}_{s,i,l,c} + \Delta q_{1,1} \hat{y}_{1,1,1,c} \geq \beta_c,\ \forall\ c \in \mathcal{C},\ K_c > 1, && \text{(A.9)} \\
& \hat{y}_{s,i,l,c} = 0,\ \forall\ c \in \mathcal{C},\ l \in \mathcal{L}_c,\ i \in \mathcal{I},\ s > K_c, && \text{(A.10)} \\
& \hat{x}_{s,i,l} \in [0,1],\ \forall\ (s,i,l) \in \mathcal{B}_{-1}, \\
& \hat{y}_{s,i,l,c} \in [0,1],\ \forall\ (s,i,l) \in \mathcal{S} \times \mathcal{I} \times \mathcal{L},\ c \in \mathcal{C}.
\end{aligned}
\right.
$$

Since $f_{i,l}(\cdot)$ is a convex function and minimization preserves convexity, function $\mathcal{F}(z)$ is a convex function of z. For a given value of the z, we can obtain the value of $\mathcal{F}(z)$ by solving the above optimization problem. Thus, we can find the optimal value of $\hat{x}_{1,1,1}$ by solving the following problem:

$$\min_{0\leq \hat{x}_{1,1,1}\leq 1} \quad \Delta q_{1,1} f_{1,1}(\hat{x}_{1,1,1}) + \mathcal{F}(\Delta q_{1,1}\hat{x}_{1,1,1}).$$

Let $\hat{x}^*_{1,1,1}(\vec{\beta})$ be an optimal solution of the above problem. Then, the first-order condition for the above optimization problem results in

$$f'_{1,1}\big(\hat{x}^*_{1,1,1}(\vec{\beta})\big) + \mathcal{F}'\big(\Delta q_{1,1}\hat{x}^*_{1,1,1}(\vec{\beta})\big) = 0. \tag{A.11}$$

Since $f_{1,1}(\cdot)$ and $\mathcal{F}(\cdot)$ are convex functions, $f'_{1,1}(\cdot)$ and $\mathcal{F}'(\cdot)$ are increasing functions. It now immediately follows that $\hat{x}^*_{1,1,1}(\vec{\beta})$ decreases with an increase in $q_{1,1}$. ■

Proof of Corollary 3: We want to find a bound for the ratio

$$\frac{\mathbb{E}[G(\bar{\mathbf{M}})]}{G(\alpha\mathbf{M})}\ . \tag{A.12}$$

Recall that $G(\bar{\mathbf{M}})$ is the objective function of problem $\mathcal{P}_1(\bar{\mathbf{M}})$ and $G(\alpha\mathbf{M})$ is the objective function of problem $\mathcal{P}_1(\alpha\mathbf{M})$. Let $\pi(\alpha\mathbf{M}) = \big(\hat{x}^*(\alpha\mathbf{M}),\ \hat{y}^*(\alpha\mathbf{M})\big)$ be the solution of problem $\mathcal{P}_1(\alpha\mathbf{M})$, where $\hat{x}^*(\alpha\mathbf{M}) = \big(\hat{x}^*_{i,l}(\alpha\mathbf{M}) :\ i\in\mathcal{I},\ l\in\mathcal{L}\big)$ and $\hat{y}^*(\alpha\mathbf{M}) = \big(\hat{y}^*_{i,l,c}(\alpha\mathbf{M}) :\ i\in\mathcal{I},\ l\in\mathcal{L},\ c\in\mathcal{C}\big)$.

We can obtain a feasible solution for problem $\mathcal{P}_1(\bar{\mathbf{M}})$ by simply multiplying the solution $\big(\hat{x}^*(\alpha\mathbf{M}),\ \hat{y}^*(\alpha\mathbf{M})\big)$ of problem $\mathcal{P}_1(\alpha\mathbf{M})$ by a factor

$$\gamma^r = \max_{c\in\mathcal{C}}\left[\frac{\bar{M}_c}{\alpha M_c}\right],$$

where r stands for our rolling-horizon policy. Since $\bar{M}_c$ is a random variable, γ^r is also a random variable. Therefore, we have

$$\frac{\mathbb{E}[G(\bar{\mathbf{M}})]}{G(\alpha\mathbf{M})} \leq \frac{\Delta\sum_{i\in\mathcal{I}}\sum_{l\in\mathcal{L}}\mathbb{E}_{\gamma^r}\Big[q_{i,l}f_{i,l}\big(\gamma^r\hat{x}^*_{i,l}(\alpha\mathbf{M})\big)\Big]}{\Delta\sum_{i\in\mathcal{I}}\sum_{l\in\mathcal{L}}\Big[q_{i,l}f_{i,l}\big(\hat{x}^*_{i,l}(\alpha\mathbf{M})\big)\Big]}\ . \tag{A.13}$$

We know that

$$\bar{M}_{s,c} = \sum_{\hat{s}=s-K+1}^{s}\frac{\tilde{M}_{\hat{s},c}}{K} + z_\alpha\sqrt{\frac{\tilde{M}_{\hat{s},c}}{K}}.$$

We know that the support of the random variable $\tilde{M}_{\hat{s},c}$ is $[0,\ M_c^{max}]$. Thus, we have $\bar{M}_{s,c} \ \leq\ M_c^{max} + z_\alpha\sqrt{KM_c^{max}}$.

We now define

$$\gamma_{max}^r = \max_{c\in\mathcal{C}}\left[\frac{M_c^{max} + z_\alpha\sqrt{KM_c^{max}}}{\alpha M_c}\right].$$

It is easy to see that $\gamma^r \leq \gamma_{max}^r$ with probability 1. Thus, from (A.13) we have

$$\frac{\mathbb{E}[G(\bar{\mathbf{M}})]}{G(\alpha\mathbf{M})} \leq \frac{\sum_{i\in\mathcal{I}}\sum_{l\in\mathcal{L}}\Big[q_{i,l}f_{i,l}\big(\gamma_{max}^r\hat{x}_{i,l}^*(\alpha\mathbf{M})\big)\Big]}{\sum_{i\in\mathcal{I}}\sum_{l\in\mathcal{L}}\Big[q_{i,l}f_{i,l}\big(\hat{x}_{i,l}^*(\alpha\mathbf{M})\big)\Big]} . \tag{A.14}$$

Now, consider the ratio,

$$\frac{f_{i,l}\big(\gamma_{max}^r\hat{x}_{i,l}^*(\alpha\mathbf{M})\big)}{f_{i,l}\big(\hat{x}_{i,l}^*(\alpha\mathbf{M})\big)}.$$

By following an analysis identical to that in the proof of Theorem 2, we obtain

$$\frac{f_{i,l}\big(\gamma_{max}^r\hat{x}_{i,l}^*(\alpha\mathbf{M})\big)}{f_{i,l}\big(\hat{x}_{i,l}^*(\alpha\mathbf{M})\big)} \leq \frac{1}{1-\frac{(\gamma_{max}^r-1)}{\gamma_{max}^r}\overline{\psi}_r}. \tag{A.15}$$

Using (A.12) and (A.15), we obtain the desired result that

$$\frac{\mathbb{E}[G(\bar{\mathbf{M}})]}{G(\alpha\mathbf{M})} \leq \frac{1}{1-\frac{(\gamma_{max}^r-1)}{\gamma_{max}^r}\overline{\psi}_r}.$$

This completes the proof of Corollary 3. ■

Details of the Regression for Estimating the Win-Curves Using Real Data (Section 3.2)

Here we provide the details of our logistic regression to estimate the win-curves. A few fields in our dataset are: (a) winning status (binary variable taking value 1 when the bid is won, 0 otherwise), (b) click status (binary variable taking value 1 if the impression is clicked, 0 otherwise), (c) bid amount, (d) zip-code, (e) device maker (Apple, Samsung, etc.), (f) operating system (iOS, android, etc.). In the data set, we have 52,460 records corresponding to 5 different zip-codes in the Boston area. We ran five different regressions, corresponding to each of the five zip-codes, to estimate the win-curves in the respective locations. The bid amount is the independent variable and the winning status is the dependent variable. After estimating the regression coefficients, we also applied the chi-square test to check the goodness-of-fit of the regression model. For location $l \in \{1,\ 2,\ 3,\ 4,\ 5\}$, (locations 1, 2, 3, 4, and 5, corresponding to zip-codes 02110, 02114, 02116, 02118, and 02119, respectively) we estimate the win-probability for a bid b as $x_l(b) = \frac{e^{\beta_0^l+\beta_1^l b}}{1+e^{\beta_0^l+\beta_1^l b}}$. Then, we obtain the expected-cost function for location l using the relation $f_l(x) = x_l b_l(x)$, where $b_l(\cdot)$ is the inverse function of $x_l(\cdot)$. The estimated function $f_l(\cdot)$ is convex at each location, as is assumed in our theoretical analysis (see Figure 3.1). The results of the regression and the goodness-of-fit test for each zip-code are in Table A.1.

Table A.1. Results of the Regression

Zip Code	**Number of Records**	**Regression Constant (β_0^l)**	**Regression Coefficient for Bid (β_1^l)**	**Chi-Square Test Statistic**
02110	1086	-2.281^{***} (std. error: 0.2141)	0.7051^{***} (std. error: 0.1290)	31.23302^{***}
02114	27693	-2.291^{***} (std. error: 0.03693)	1.04294^{***} (std. error: 0.02224)	2415.644^{***}
02116	6794	-1.905^{***} (std. error: 0.07184)	0.87651^{***} (std. error: 0.04388)	426.4403^{***}
02118	2833	-1.936^{***} (std. error: 0.11432)	0.76762^{***} (std. error: 0.06958)	128.1244^{***}
02119	14054	-2.193^{***} (std. error: 0.05212)	0.97615^{***} (std. error: 0.03133)	1053.156^{***}

Significance codes: '***' $p < 0.001$, '**' $p < 0.01$, '*' $p < 0.05$, '.' $p < 0.1$. The chi-square test with p-value of less than 0.001 tells us that our model as a whole fits significantly better than an empty model.

Proof of Theorem 4: Recall from Section 4.2.1 that problem P_{After} is defined as follows:

$$\max_{I_G, a_b, a_r} \quad R(I_G, a_b, a_r) \quad \text{s.t. } I_G \in \{0,1\} \text{ and } a_r, a_b \in \{0\} \cup [a_{\min}, \infty).$$

- **Overview of the Solution Procedure:** For a given value of the gating decision I_G, we first ignore the constraints that a_b and a_r should be at least $a_{\min}$ and obtain the unconstrained optimal values of the ad-intensities. Let a_b^u and a_r^u denote these unconstrained optimal ad-intensities for ad-blockers and regulars, respectively. Then, the consideration of $a_{\min}$ leads to the following three possibilities: (i) $a_{\min} < \min\{a_b^u, a_r^u\}$, (ii) $a_{\min} > \max\{a_b^u, a_r^u\}$, and (iii) $\min\{a_b^u, a_r^u\} \leq a_{\min} \leq \max\{a_b^u, a_r^u\}$. We separately solve for the optimal decisions in each of these three ranges. As we will see below, the three ranges of $a_{\min}$ correspond to the three cases in the statement of Theorem 4.

We now proceed with the analysis. Recall from (4.2) that the equilibrium revenue of the website is

$$R(I_G, a_b, a_r) = r\Big[n_b(I_G, a_b, a_r) \cdot a_b + n_r(I_G, a_b, a_r) \cdot a_r\Big].$$

We will obtain the equilibrium expressions for $n_b(I_G, a_b, a_r)$ and $n_r(I_G, a_b, a_r)$. Using the same argument as in the proof of Property A3, we see that it is sufficient to consider only the set of feasible decisions (I_G, a_b, a_r) that lead to an equilibrium value of v satisfying the following conditions:

$$0 \leq v - u_0 \leq C_r a_r,\ 0 \leq v - u_0 \leq C_b a_b. \tag{A.16}$$

If the website gates ad-blockers, then the probability that a randomly chosen ad-blocker will white-list the website is $\frac{v-u_0}{C_b a_b}$. On the other hand, if the website does not gate ad-blockers, then all the ad-blockers become users of the website (because $v - u_0 \geq 0$). Thus, we have

$$n_b(I_G, a_b, a_r) = \frac{NBI_G(v - u_0)}{C_b a_b} + (1 - I_G)NB. \tag{A.17}$$

Similarly, the probability that a randomly chosen regular will become a user of the website is $\frac{v-u_0}{C_r a_r}$. Thus, we have

$$n_r(I_G, a_b, a_r) = \frac{N(1 - B)(v - u_0)}{C_r a_r}. \tag{A.18}$$

The equilibrium traffic of the website can now be written as

$$\begin{aligned} n(I_G, a_b, a_r) &= n_b(I_G, a_b, a_r) + n_r(I_G, a_b, a_r), \\ &= \frac{NBI_G(v-u_0)}{C_b a_b} + (1-I_G)NB + \frac{N(1-B)(v-u_0)}{C_r a_r}. \end{aligned} \tag{A.19}$$

Let $A_1(I_G) = V\bar{\alpha} + V\alpha(1-I_G)B$ and $A_2(a_b, a_r) = V\alpha\left[\frac{I_G B}{C_b a_b} + \frac{1-B}{C_r a_r}\right]$. Using (4.1) and (A.19), we get

$$v = \frac{A_1(I_G) - u_0 A_2(a_b, a_r)}{1 - A_2(a_b, a_r)}. \tag{A.20}$$

Recall from Section 4.2.1 that

$$\phi(I_G, a_b, a_r) = v - u_0 = \frac{V\bar{\alpha} + V\alpha(1-I_G)B - u_0}{1 - V\alpha\left(\frac{I_G B}{C_b a_b} + \frac{1-B}{C_r a_r}\right)}.$$

Thus, from (A.16) and the above equation, we have

$$\phi(I_G, a_b, a_r) \geq 0, \tag{A.21}$$

$$a_r \geq \frac{V - V\alpha I_G B - u_0}{C_r\left[1 - \frac{V\alpha I_G B}{C_b a_b}\right]}, \tag{A.22}$$

$$a_b \geq \frac{V\bar{\alpha} + V\alpha B - u_0}{C_b\left[1 - \frac{V\alpha(1-B)}{C_r a_r}\right]}. \tag{A.23}$$

Substituting the expression of v from (A.20) in (A.17) and (A.18), we obtain expressions for $n_b(I_G, a_b, a_r)$ and $n_r(I_G, a_b, a_r)$ in terms of our three decision variables. Using these expressions in (4.2), we have

$$R(I_G, a_b, a_r) = \frac{Nr[A_1(I_G) - u_0]}{1 - A_2(a_b, a_r)}\left[\frac{BI_G}{C_b} + \frac{1-B}{C_r}\right]. \tag{A.24}$$

Since $A_2(a_b, a_r)$ is decreasing in both a_b and a_r, it is easy to see that $R(I_G, a_b, a_r)$ is also decreasing in a_b and a_r. Therefore, the unconstrained optimal values of a_b and a_r are the lowest feasible values. From (A.22) and (A.23), we have

$$a_r^u = \frac{V - V\alpha I_G B - u_0}{C_r\left[1 - \frac{V\alpha I_G B}{C_b a_b^u}\right]}, \tag{A.25}$$

$$a_b^u = \frac{V\bar{\alpha} + V\alpha B - u_0}{C_b\left[1 - \frac{V\alpha(1-B)}{C_r a_r^u}\right]}. \tag{A.26}$$

Solving for a_b^u and a_r^u, we have

$$a_b^u = \frac{V - u_0}{C_b},\ a_r^u = \frac{V - u_0}{C_r}. \tag{A.27}$$

Having obtained the unconstrained ad-intensities, we now consider the constraints that $a_b, a_r \geq a_{\min}$. Since $C_b \geq C_r$, we have $a_b^u \leq a_r^u$. Thus, the three cases mentioned in the overview of this proof translate into: (i) $a_{\min} < a_b^u$, (ii) $a_{\min} > a_r^u$, and (iii) $a_b^u \leq a_{\min} \leq a_r^u$. We now proceed to obtain the optimal gating decision in each of the three cases.

Case (i) ($a_{\min} < a_b^u$): Using (A.27) and $C_b \geq C_r$, it is easy to see that the condition $a_{\min} < a_b^u$ is identical to $C_b,\ C_r < \widehat{C}$. In other words, this case corresponds to the class of websites with low C_b and low C_r. Since $a_{\min} < a_b^u$, we have $a_b^u,\ a_r^u > a_{\min}$. Thus, the unconstrained optimal solution in (A.27) is feasible for the constrained problem, and, therefore, we have $a_b^* = a_b^u$ and $a_r^* = a_r^u$. Substituting a_b^* and a_r^* in (A.24), we get

$$R(I_G, a_b^*, a_r^*) = N(V - u_0)r\left[\frac{BI_G}{C_b} + \frac{1-B}{C_r}\right].$$

It is easy to see that $R(1, a_b^*, a_r^*) \geq R(0, a_b^*, a_r^*)$; therefore, $I_G^* = 1$. The optimal solution is:

$$(I_G^*,\ a_b^*,\ a_r^*) = \left(1,\ \frac{V - u_0}{C_b},\ \frac{V - u_0}{C_r}\right).$$

Case (ii) ($a_{\min} > a_r^u$): Using (A.27) and $C_b \geq C_r$, it is easy to see that the condition $a_{\min} > a_r^u$ is identical to $C_b,\ C_r > \widehat{C}$. Thus, this case corresponds to the class of websites with high

C_b and high C_r. Here, the unconstrained optimal values of a_b and a_r are both lower than $a_{\min}$. Since $R(I_G, a_b, a_r)$ is decreasing in a_b and a_r, we have $a_b^* = a_r^* = a_{\min}$. Substituting a_b^* and a_r^* in (A.24), we get

$$R(I_G, a_b^*, a_r^*) = \frac{Nr[A_1(I_G) - u_0]}{1 - A_2(a_{\min}, a_{\min})} \left[\frac{BI_G}{C_b} + \frac{1-B}{C_r}\right].$$

Comparing the revenues at $I_G = 0$ and $I_G = 1$, we obtain the following optimal solution:

$$(I_G^*,\ a_b^*,\ a_r^*) = \begin{cases} (0,\ a_{\min},\ a_{\min}), & \text{if } u_0 > u_h,\ \phi(0,\ a_{\min},\ a_{\min}) \geq 0, \\ (1,\ a_{\min},\ a_{\min}), & \text{if } u_0 \leq u_h,\ \phi(1,\ a_{\min},\ a_{\min}) \geq 0, \\ (0,\ 0,\ 0), & \text{otherwise.} \end{cases}$$

Case (iii) ($a_b^u \leq a_{\min} \leq a_r^u$): Using (A.27) and $C_b \geq C_r$, it is easy to see that the condition $a_b^u \leq a_{\min} \leq a_r^u$ is identical to $C_r \leq \widehat{C} \leq C_b$. Thus, this case corresponds to the class of websites with high C_b and low C_r. Since a_b^u is lower than $a_{\min}$ in this case, and $R(I_G, a_b, a_r)$ is decreasing in a_b, the optimal value of a_b is $a_{\min}$; that is, $a_b^* = a_{\min}$. Thus, from (A.25), we have

$$a_r^* = \frac{V - V\alpha I_G B - u_0}{C_r\left[1 - \frac{V\alpha I_G B}{C_b a_{\min}}\right]}.$$

Substituting a_b^* and a_r^* in (A.24) we get

$$R(I_G, a_b^*, a_r^*) = Nr\left[\frac{V - BI_G V\alpha - u_0}{1 - \frac{V\alpha BI_G}{a_{\min} C_b}}\right]\left[\frac{BI_G}{C_b} + \frac{1-B}{C_r}\right].$$

Comparing the revenue at $I_G = 0$ and $I_G = 1$, we obtain the following optimal solution:

$$(I_G^*,\ a_b^*,\ a_r^*) = \begin{cases} (0,\ a_{\min},\ \frac{V-u_0}{C_r}), & \text{if } u_0 > u_m,\ \phi(0,\ a_{\min},\ \frac{V-u_0}{C_r}) \geq 0, \\ (1,\ a_{\min},\ a_g) & \text{if } u_0 \leq u_m,\ \phi(1,\ a_{\min},\ a_g) \geq 0, \\ (0,\ 0,\ 0), & \text{otherwise.} \end{cases}$$

This completes the proof of Theorem 4. ■

Proof of Theorem 5: From our analysis in the proof of Theorem 4, we also obtain the optimal ad-intensities for a *given* value of the gating decision I_G. We note these below:

- **Case A:** If $I_G = 0$, then
  - **Case A1:** $(\hat{a}_b, \hat{a}_r) = (a_{\min}, \frac{V-u_0}{C_r}),$ if $a_{\min} \leq \frac{V-u_0}{C_r}, \;\; \phi(0, a_{\min}, \frac{V-u_0}{C_r}) \geq 0$,
  - **Case A2:** $(\hat{a}_b, \hat{a}_r) = (a_{\min}, a_{\min}),$ if $a_{\min} > \frac{V-u_0}{C_r}, \;\; \phi(0, a_{\min}, a_{\min}) \geq 0,$
  - **Case A3:** $(\hat{a}_b, \hat{a}_r) = (0, 0),$ otherwise.

- **Case B:** If $I_G = 1$, then
  - **Case B1:** $(\hat{a}_b, \hat{a}_r) = (\frac{V-u_0}{C_b}, \frac{V-u_0}{C_r}),$ if $a_{\min} < \frac{V-u_0}{C_b}, \;\; \phi(1, \frac{V-u_0}{C_b}, \frac{V-u_0}{C_r}) \geq 0,$
  - **Case B2:** $(\hat{a}_b, \hat{a}_r) = (a_{\min}, a_g),$ if $\frac{V-u_0}{C_b} \leq a_{\min} \leq \frac{V-u_0}{C_r}, \;\; \phi(1, a_{\min}, a_g) \geq 0,$
  - **Case B3:** $(\hat{a}_b, \hat{a}_r) = (a_{\min}, a_{\min}),$ if $a_{\min} > \frac{V-u_0}{C_r}, \;\; \phi(1, a_{\min}, a_{\min}) \geq 0,$
  - **Case B4:** $(\hat{a}_b, \hat{a}_r) = (0, 0),$ otherwise.

Consider Case A (i.e., $I_G = 0$). Then, we have the following:

- If the optimal ad-intensities $(\hat{a}_b, \hat{a}_r)$ fall under Case A1, then an increase in V keeps us in Case A1. This is because (a) the condition $a_{\min} \leq \frac{V-u_0}{C_r}$ continues to hold and (b) the condition $\phi(0, a_{\min}, \frac{V-u_0}{C_r}) \geq 0$ also continues to hold since the left-hand-side is an increasing function of V.

- If the optimal ad-intensities $(\hat{a}_b, \hat{a}_r)$ fall under Case A2, then an increase in V either keeps us in Case A2, in which case the result holds trivially, or we switch to Case A1, where both the ad-intensities are at least as large as their corresponding values under Case A2 (since $\frac{V-u_0}{C_r} \geq a_{\min}$).

- If the optimal ad-intensities fall under Case A3, then the result holds trivially.

The argument is similar for Case B ($I_G = 1$).

To show that the optimal gating decision I_G^* is not necessarily monotone in V, it suffices to exhibit a numerical counter-example. To this end, let $C_r = 100$, $C_b = 200$, $a_{\min} = 2$, $N = 10^6$, $r = 1$, $\alpha = 0.99$, $B = 0.01$, and $u_0 = 50$. It is easy to verify using Theorem 4 that, for these parameters, $I_G^* = 1$ at $V = 240$, and $I_G^* = 0$ at $V = 260$. ∎

Proof of Theorem 6: Recall from Section 4.2.2 that problem P_{Before} is as follows:

$$\max_a \quad R(1, a, a) \quad \text{s.t. } a \in \{0\} \cup [a_{\min}, \infty).$$

Similar to our analysis of problem P_{After}, here too we will first ignore the constraint $a \geq a_{\min}$ and solve the relaxed problem. Let a^u denote the optimal ad-intensity in this unconstrained problem. Since $n(1, a, a)$ is the equilibrium traffic of the website, its revenue is

$$R(1, a, a) = n(1, a, a)ra. \tag{A.28}$$

We now proceed to obtain the optimal ad-intensity of the relaxed problem. Since $C_b \geq C_r$, using property B2 in Section 4.2.2, we have the following two possibilities depending on the equilibrium value v:

- **Case (i) ($0 \leq v - u_0 \leq aC_r$):** This case corresponds to the value premium offered by the website over the outside option, $v - u_0$, being low. Since $C_b \geq C_r$, the probability that a randomly chosen low-cost potential user becomes the user of the website is $\frac{v-u_0}{aC_r}$. Similarly, the probability that a randomly chosen high-cost potential user becomes the user of the website is $\frac{v-u_0}{aC_b}$. Thus, in this case, the equilibrium traffic of the website is

$$n(1, a, a) = \frac{N(1-B)(v-u_0)}{aC_r} + \frac{NB(v-u_0)}{aC_b}. \tag{A.29}$$

Let $A_3(a) = \frac{V\alpha(1-\rho)}{aC_r}$. Recall from Section 4.2.2 the two parametric constants $\hat{a} = \frac{V(1-\alpha\rho)-u_0}{C_r}$ and $\hat{\hat{a}} = \frac{V-u_0}{C_b}$. Using (4.1) and (A.29), we obtain

$$v = \frac{V\bar{\alpha} - u_0 A_3(a)}{1 - A_3(a)}. \tag{A.30}$$

Using $0 \le v - u_0 \le aC_r$ and (A.30), we have

$$\phi_l(a) \ge 0, \; a \ge \hat{a}. \tag{A.31}$$

Substituting the expression of v from (A.30) in (A.29) leads to the expression for $n(1,a,a)$ in terms of the decision variable a. Using this expression in (A.28), we have

$$R(1,a,a) = \frac{Nr(1-\rho)}{C_r}\left[\frac{V\bar{\alpha} - u_0}{1 - A_3(a)}\right]. \tag{A.32}$$

Since $A_3(a)$ is decreasing in a, $R(1,a,a)$ is decreasing in a. Thus, the optimal value of a is its minimum feasible value. Using (A.31), we therefore have $a^u = \hat{a}$.

- **Case (ii) ($aC_r \le v - u_0 \le aC_b$):** Here, the value premium offered by the website is high. Since $v - u_0 \ge aC_r$, all the low-cost potential users become users of the website. However, the probability that a randomly chosen high-cost potential user becomes a user of the website is $\frac{v-u_0}{aC_b}$. Thus, the equilibrium traffic of the website is

$$n(1,a,a) = N(1-B) + \frac{NB(v-u_0)}{aC_b}. \tag{A.33}$$

Let $A_4(a) = \frac{V\alpha B}{aC_b}$. Using (4.1) and (A.33), we obtain

$$v = \frac{(V\bar{\alpha} + V\alpha\bar{B}) - u_0 A_4(a)}{1 - A_4(a)}. \tag{A.34}$$

Using $aC_r \le v - u_0 \le aC_b$ and (A.34), we have

$$\phi_h(a) \ge 0, \; \hat{\hat{a}} \le a \le \hat{a}. \tag{A.35}$$

Substituting the expression of v from (A.34) in (A.33), we obtain the expression of $n(1,a,a)$ in terms of the decision variable a. Using this expression in (A.28), we have

$$R(1,a,a) = \left[\bar{B} + \frac{B(V\bar{\alpha} + V\alpha\bar{B} - u_0)}{aC_b - V\alpha B}\right] Nra. \tag{A.36}$$

Thus,

$$\frac{\partial^2 R(1,a,a)}{\partial a^2} = \frac{2V\alpha^2 B^3 Nr}{C_b^3 a^4} \frac{[V\bar{\alpha} + V\alpha\bar{B} - u_0]}{(1 - V\alpha\frac{B}{aC_b})^3}. \tag{A.37}$$

Using (A.35), it is straightforward to show that $1 - V\alpha\frac{B}{aC_b} \geq 0$ and $u_0 \leq V\bar{\alpha} + V\alpha\bar{B}$. Thus, we have $\frac{\partial^2 R(1,a,a)}{\partial a^2} \geq 0$. Thus, $R(1,a,a)$ is a convex function of a, and the optimal value of a is at the boundary of its feasible region. Using (A.35) and observing that $R(1,\hat{\hat{a}},\hat{\hat{a}}) \leq R(1,\hat{a},\hat{a})$ if and only if $u_0 \leq \hat{u}_0$, we have $a^u = \hat{a}$ when $u_0 \leq \hat{u}_0$; otherwise $a^u = \hat{\hat{a}}$.

Combining the above two cases, we have

$$a^u = \begin{cases} \hat{\hat{a}}, \text{ if } u_0 > \hat{u}_0, \\ \hat{a}, \text{ if } u_0 \leq \hat{u}_0. \end{cases} \tag{A.38}$$

We now consider the constraint $a \geq a_{\min}$. Since $\hat{\hat{a}} \leq \hat{a}$, we have the following three cases: (a) $a_{\min} < \hat{\hat{a}}$, (b) $a_{\min} > \hat{a}$, and (c) $\hat{\hat{a}} \leq a_{\min} \leq \hat{a}$. These correspond to the three cases in the statement of Theorem 6.

Case (a) ($a_{\min} < \hat{\hat{a}}$): Using (A.38) and $C_b \geq C_r$, the condition $a_{\min} < \hat{\hat{a}}$ is equivalent to $C_b, C_r < \widehat{C}$. Accordingly, this case corresponds to the class of websites with low C_b and low C_r. Since $a_{\min} < \hat{\hat{a}}$, we have $a^u > a_{\min}$. Thus, the unconstrained optimal solution in (A.38) is feasible to problem P_{Before}. The optimal solution in this case is:

$$a^* = \begin{cases} \hat{a}, & \text{if } u_0 \leq \hat{u}_0,\ \phi_h(\hat{a}) \geq 0, \\ \hat{\hat{a}}, & \text{if } u_0 > \hat{u}_0,\ \phi_h(\hat{\hat{a}}) \geq 0, \\ 0, & \text{otherwise}. \end{cases}$$

Case (b) ($a_{\min} > \hat{a}$): Using (A.38) and $C_b \geq C_r$, the condition $a_{\min} > \hat{a}$ is identical to $C_b \geq \widehat{C},\ C_r > \widehat{C}_r$. Therefore, this case corresponds to the class of websites with high C_b and high C_r. In this case, the unconstrained optimal value of a is lower than $a_{\min}$. Thus, the optimal solution in this case is:

$$a^* = \begin{cases} a_{\min}, & \text{if } \phi_l(a_{\min}) \geq 0, \\ 0, & \text{otherwise.} \end{cases}$$

Case (c) ($\hat{\hat{a}} \leq a_{\min} \leq \hat{a}$): Using (A.38) and $C_b \geq C_r$, it is easy to see that the condition $\hat{\hat{a}} \leq a_{\min} \leq \hat{a}$ is identical to $C_b \geq \widehat{C},\ C_r \leq \widehat{C}_r$. Thus, this case corresponds to the class of websites with high C_b and low C_r. The optimal ad-intensity in this case is:

$$a^* = \begin{cases} \hat{a}, & \text{if } u_0 \leq u_b,\ \phi_h(\hat{a}) \geq 0, \\ a_{\min}, & \text{if } u_0 > u_b,\ \phi_h(a_{\min}) \geq 0, \\ 0, & \text{otherwise.} \end{cases}$$

This completes the proof of Theorem 6. ■

Proof of Theorem 8: Let a^*_{Before} be an optimal ad-intensity in the pre-ad-block world. Using the argument presented immediately after the statement of the theorem in Section 4.3, we have that the website's revenue in the post-ad-block world under the feasible solution $(I_G, a_b, a_r) = (1, a^*_{Before}, a^*_{Before})$ is exactly the same as its optimal revenue in the pre-ad-block world. Thus, the revenue of the website in the pre-ad-block world is a lower bound on its revenue in the post-ad-block world. The result follows. ■

Proof of Theorem 9: Using the optimal solutions in the post- and pre-ad-block worlds from Theorems 4 and Theorem 6, respectively, we summarize the optimal ad-intensities in Tables A.2 and A.3 below. The notation used here is as defined in Sections 4.2.1 and 4.2.2.

Table A.2. The optimal ad-intensities in the pre- and post-ad-block worlds, when the outside option $u_0 > u_m$.

$u_0 > u_m$	Pre-Ad-Block	Post-Ad-Block
$a_{\min} \leq \frac{V-u_0}{C_b}$	$a^* = \begin{cases} \hat{a} = \frac{V-V\alpha\rho-u_0}{C_r}, & \text{if } u_0 \leq \hat{u}_0, \\ \hat{\hat{a}} = \frac{V-u_0}{C_b}, & \text{otherwise.} \end{cases}$	$a_b^* = \frac{V-u_0}{C_b}$ $a_r^* = \frac{V-u_0}{C_r}$
$\frac{V-u_0}{C_b} \leq a_{\min} \leq \frac{V-u_0}{C_r}$	$a^* = \begin{cases} \hat{a}, & \text{if } a_{\min} \leq \hat{a},\ u_0 \leq \hat{u}_0, \\ a_{\min}, & \text{otherwise.} \end{cases}$	$a_b^* = 0$ $a_r^* = \frac{V-u_0}{C_r}$
$a_{\min} \geq \frac{V-u_0}{C_r}$	$a^* = a_{\min}$	$a_b^* = a_r^* = a_{\min}$

Table A.3. The optimal ad-intensities in the pre- and post-ad-block worlds, when the outside option $u_0 \leq u_m$.

$u_0 \leq u_m$	Pre-Ad-Block	Post-Ad-Block
$a_{\min} \leq \frac{V-u_0}{C_b}$	$a^* = \begin{cases} \hat{a} = \frac{V-V\alpha\rho-u_0}{C_r}, & \text{if } u_0 \leq \hat{u}_0, \\ \hat{\hat{a}} = \frac{V-u_0}{C_b}, & \text{otherwise.} \end{cases}$	$a_b^* = \frac{V-u_0}{C_b}$ $a_r^* = \frac{V-u_0}{C_r}$
$\frac{V-u_0}{C_b} \leq a_{\min} \leq a_g$	$a^* = \begin{cases} \hat{a} = \frac{V-V\alpha\rho-u_0}{C_r}, & \text{if } a_{\min} \leq \hat{a},\ u_0 \leq \hat{u}_0, \\ a_{\min}, & \text{otherwise.} \end{cases}$	$a_b^* = a_{\min}$ $a_r^* = a_g$
$a_{\min} \geq a_g$	$a^* = a_{\min}$	$a_b^* = a_r^* = a_{\min}$

In each of the cases in Tables A.2 and A.3, it is straightforward to verify that $a_b^* \leq a^*$ and $a_r^* \geq a^*$. ■

Proof of Theorem 10: From the proof of Theorem 4, when $a_{\min} \leq \frac{V-u_0}{C_b}$, we have the following for Problem P_{After} (i) $n = N$, (ii) $v = V$, and (iii) $a_b^* = \frac{V-u_0}{C_b}$, and (iv) $a_r^* = \frac{V-u_0}{C_r}$. Thus, we have

$$CS_{After} = N(1-B)\left(V - \frac{a_r^* C_r}{2}\right) + NB\left(V - \frac{a_b^* C_b}{2}\right).$$

Using $a_r^* = \frac{V-u_0}{C_r}$ and $a_b^* = \frac{V-u_0}{C_b}$, we get

$$CS_{After} = \frac{N(V+u_0)}{2}. \tag{A.39}$$

For Problem P_{Before}, from the proof of Theorem 6, the optimal ad-intensity when $a_{\min} \leq \frac{V-u_0}{C_b}$ is:

$$a^* = \begin{cases} \hat{a} = \frac{V - V\alpha\rho - u_0}{C_r}, & \text{if } u_0 \leq \hat{u}_0, \\ \\ \hat{\hat{a}} = \frac{V-u_0}{C_b}, & \text{otherwise.} \end{cases}$$

We now compute the consumer surplus in the pre-ad-block world in the following two cases.

Case (i): $u_0 \leq \hat{u}_0$. In this case, we have $a^* = \hat{a}$. Let $v(\hat{a})$ be the value of the website when the ad-intensity is $\hat{a}$. Then, the consumer surplus can be written as

$$CS_{Before} = N\bar{B}\left[v(\hat{a}) - \frac{\hat{a}C_r}{2}\right] + \frac{NB[v(\hat{a}) - u_0]}{\hat{a}C_b}\left[v(\hat{a}) - \frac{v(\hat{a}) - u_0}{2}\right] + NB\left[1 - \frac{v(\hat{a}) - u_0}{\hat{a}C_b}\right]u_0. \tag{A.40}$$

Using $v(\hat{a}) - u_0 = \hat{a}C_r$ and $\hat{a} = \frac{V - V\alpha\rho - u_0}{C_r}$, we get

$$CS_{Before} = \frac{N(1-\rho)(V - V\alpha\rho + u_0)}{2} + NB\rho u_0.$$

Using the definition of $\hat{u}_0$, it can be easily verified that $\hat{u}_0 \leq \frac{V[1-(1-\rho)(1-\alpha\rho)]}{\rho(2B-1)}$. Since $u_0 \leq \hat{u}_0$, we have $u_0 \leq \frac{V[1-(1-\rho)(1-\alpha\rho)]}{\rho(2B-1)}$, which implies that $CS_{After} \geq CS_{Before}$.

Case (ii): $u_0 > \hat{u}_0$. In this case, we know that $a^* = \hat{\hat{a}}$. Then, from the proof of Theorem 6, we have $n = N$ and $v = V$. Thus,

$$CS_{Before} = N(1-B)\left[V - \frac{\hat{\hat{a}}C_r}{2}\right] + NB\left[V - \frac{\hat{\hat{a}}C_b}{2}\right].$$

Using $\hat{\hat{a}} = \frac{V-u_0}{C_b}$, we get

$$CS_{Before} = N(1-B)\left[V - \frac{(V-u_0)}{C_b}\frac{C_r}{2}\right] + \frac{NB(V+u_0)}{2}.$$

Clearly, CS_{Before} is increasing in C_b. Further, whenever $C_b = C_r$, we have $CS_{Before} = CS_{After}$. Thus, for $C_b \geq C_r$, we have $CS_{After} \leq CS_{Before}$.

To summarize, consumer surplus increases if $u_0 \leq \hat{u}_0$ and decreases otherwise. Further, it is easy to verify that the above arguments are reversible; i.e., the condition

$$CS_{After} \geq CS_{Before}$$

implies $u_0 \leq \hat{u}_0$. ■

Proof of Theorem 11: From Theorem 8, we know that, for any u_0, the revenue of the website increases in the post-ad-block world. Further, from Theorem 10, consumer surplus increases in the post-ad-block world if $u_0 \leq \hat{u}_0$, and decreases otherwise. Thus, if $u_0 \leq \hat{u}_0$, both consumer surplus and the website's revenue increase in the post-ad-block world and hence the social surplus also increases. If $u_0 > \hat{u}_0$, consumer surplus decreases but the revenue of the website increases. In this case, we know that (i) $a^* = \hat{\hat{a}}$, from Theorem 6, and (ii) $(I_G^*, a_b^*, a_r^*) = \left(1, \frac{V-u_0}{C_b}, \frac{V-u_0}{C_r}\right)$, from Theorem 4. Therefore,

$$W_{Before} = N(1-B)V - \frac{N(1-B)(V-u_0)C_r}{2C_b} + \frac{NB(V+u_0)}{2} + \frac{Nr(V-u_0)}{C_b},$$

and

$$W_{After} = Nr(V-u_0)\left[\frac{B}{C_b} + \frac{1-B}{C_r}\right] + \frac{N(V+u_0)}{2}.$$

It is straightforward to verify that $W_{Before} > W_{After}$. Thus, the social surplus decreases in the post-ad-block world if $u_0 > \hat{u}_0$. The result follows. ■

Proof of Theorem 12: Let $v_r = u_0 / \left[1 - \frac{\alpha B \bar{\rho}}{1-B}\right]$. Recall from Section 4.2.2 that $\rho = \frac{B(C_b - C_r)}{C_b}$. Using (4.6), (4.7), and the definitions of v_r and $\hat{u}_0$, we have

$$\frac{\partial R_{Before}}{\partial V} = \begin{cases} \frac{Nr(1-\rho)(1-\alpha\rho)}{C_r}, & \text{if } V \geq v_r,\ v_r \geq 0, \\ \\ \frac{Nr}{C_b}, & \text{otherwise,} \end{cases}$$

and

$$\frac{\partial R_{After}}{\partial V} = \frac{Nr(1-\rho)}{C_r}, \\ = Nr\left[\frac{B}{C_b} + \frac{1-B}{C_r}\right].$$

Thus,

$$\frac{\partial R_{After}}{\partial V} - \frac{\partial R_{Before}}{\partial V} = \begin{cases} \frac{Nr(1-\rho)\alpha\rho}{C_r}, & \text{if } V \geq v_r,\ v_r \geq 0, \\ \\ \frac{Nr(1-B)}{C_r}, & \text{otherwise.} \end{cases} \tag{A.41}$$

Thus, $\frac{\partial R_{After}}{\partial V} - \frac{\partial R_{Before}}{\partial V} \geq 0$. Further, it is easy to verify that equality holds at $B = 0$ and $B = 1$. ■

Proof of Theorem 13: We only need to show that $R(p, a_{\min}, a_{\min})$ is a concave function of p. To do this, we first derive the expression for this function. Using property C1 in Section 4.5, we know that $0 \leq v - u_0 \leq a_{\min}C_r$. Thus, a $\frac{v-u_0}{a_{\min}C_b}$ fraction of ad-blockers and a $\frac{v-u_0}{a_{\min}C_r}$ fraction of regulars become users of the website. Therefore, the total traffic of the website is

$$n(p, a_{\min}, a_{\min}) = \frac{NB(v-u_0)p}{a_{\min}C_b} + (1-p)NB + N(1-B)\frac{v-u_0}{a_{\min}C_r}. \tag{A.42}$$

Using (4.1) and (A.42), we have

$$v(p) = \frac{H_1 - V\alpha u_0 H_2}{1 - V\alpha H_2}, \tag{A.43}$$

where $H_1 = V\bar{\alpha} + V\alpha(1-p)B$ and $H_2 = \frac{Bp}{a_{\min}C_b} + \frac{1-B}{a_{\min}C_r}$. Thus,

$$\frac{\partial v}{\partial p} = \frac{BV\alpha}{1 - V\alpha H_2}\left[\frac{v-u_0}{a_{\min}C_b} - 1\right]. \tag{A.44}$$

Since $v - u_0 \leq a_{\min}C_b$, we have $\frac{\partial v}{\partial p} \leq 0$. Similarly, we have

$$\frac{\partial^2 v}{\partial p^2} = \frac{BV\alpha}{a_{\min}C_b(1 - V\alpha H_2)}\frac{\partial v}{\partial p} + \left(\frac{v-u_0}{a_{\min}C_b} - 1\right)\frac{Bv_0^2 V\alpha^2}{(1 - V\alpha H_2)^2}\frac{B}{a_{\min}C_b}. \tag{A.45}$$

Since $\frac{\partial v}{\partial p} \leq 0$ and $v - u_0 \leq a_{\min} C_b$, we have $\frac{\partial^2 v}{\partial p^2} \leq 0$.

The revenue of the website is

$$R(p, a_{\min}, a_{\min}) = \frac{NB(v - u_0)p}{a_{\min} C_b} r a_{\min} + N(1 - B)\frac{v - u_0}{a_{\min} C_r} r a_{\min},$$
$$= Nr(v - u_0)\left[\frac{Bp}{C_b} + \frac{1 - B}{C_r}\right].$$

Thus, we have

$$\frac{\partial R(p, a_{\min}, a_{\min})}{\partial p} = Nr(v - u_0)\frac{B}{C_b} + Nr\left[\frac{Bp}{C_b} + \frac{1 - B}{C_r}\right]\frac{\partial v}{\partial p},$$
$$\frac{\partial^2 R(p, a_{\min}, a_{\min})}{\partial p^2} = \frac{NrB}{C_b}\frac{\partial v}{\partial p} + Nr\left[\frac{Bp}{C_b} + \frac{1 - B}{C_r}\right]\frac{\partial^2 v}{\partial p^2} + \frac{\partial v}{\partial p}\frac{NrB}{C_b}.$$

Using $\frac{\partial v}{\partial p} \leq 0$ and $\frac{\partial^2 v}{\partial p^2} \leq 0$, we have $\frac{d^2 R(p, a_{\min}, a_{\min})}{\partial p^2} \leq 0$. Thus, $R(p, a_{\min}, a_{\min})$ is concave function of p. The result follows. ■

Analysis of Extensions in Section 4.6:

Here, we analyze the four extensions that were defined in Section 4.6.

Heterogeneous Profitability:

Recall the setting of this extension from Section 4.6.1. In the pre-ad-block world, when the website was unable to discriminate between users, the website's expected revenue corresponding to an ad-intensity of a was

$$[n_a(a) \times a r_a + n_g(a) \times a r_g],$$

where $n_a(a)$ is the equilibrium number of users of the website from the potential ad-block (resp., regular) user population.

It is easy to see that in the post-ad-block world, if the website gates all the ad-block users and offers the same ad-intensity a (as that in the pre-ad-block world) to all users, then the website earns the same revenue. This is because, under this specific post-ad-block strategy, the utility of every regular user (resp., every ad-block user) is the same as her utility in the pre-ad-block strategy. Thus, the number of regular users visiting the website, the number of ad-block users visiting the website, and the website's revenue under this post-ad-block strategy are all equal to their corresponding values under the pre-ad-block strategy. This observation immediately implies that the website's optimal post-ad-block revenue is greater than or equal to its optimal pre-ad-block revenue.

We now proceed to show that our other main conclusions also continue to hold. Under heterogeneous profitabilities for ad-block users and regular users, we obtain the optimal ad-intensities and the corresponding social welfare in both the pre- and post-ad-block worlds. This analysis is summarized below:

Post-Ad-Block Analysis: The analysis in the proof of Theorem 4 remains the same until (and including) equation (A.24). The revenue of the website can now be written as

$$R(I_G, a_b, a_r) = \frac{N[A_1(I_G) - u_0]}{1 - A_2(a_b, a_r)} \left[\frac{BI_G r_b}{C_b} + \frac{(1-B)r_g}{C_r} \right], \tag{A.46}$$

where $A_1(I_G)$ and $A_2(a_b, a_r)$ are as defined in the proof of Theorem 4. Using the same argument as in that proof, we have

$$a_b^* = \frac{V - u_0}{C_b},\ a_r^* = \frac{V - u_0}{C_r}, \text{ and} \tag{A.47}$$

$$R(I_G, a_b^*, a_r^*) = N(V - u_0) \left[\frac{BI_G r_b}{C_b} + \frac{(1-B)r_g}{C_r} \right].$$

It is easy to see that $R(1, a_b^*, a_r^*) \geq R(0, a_b^*, a_r^*)$; therefore, $I_G^* = 1$. Thus, we have

$$(I_G^*,\ a_b^*,\ a_r^*) = \left(1,\ \frac{V - u_0}{C_b},\ \frac{V - u_0}{C_r}\right).$$

The consumer surplus CS_{After} in the post-ad-block world is

$$CS_{After} = \frac{N(V + u_0)}{2}. \tag{A.48}$$

Thus, the social welfare W_{After} in the post-ad-block world is

$$W_{After} = N(V - u_0)\left[\frac{Br_b}{C_b} + \frac{(1-B)r_g}{C_r}\right] + \frac{N(V+u_0)}{2}.$$

Pre-Ad-block Analysis: The analysis of the pre-ad-block world continues to hold, and we have

$$a^* = \begin{cases} \hat{a}, & \text{if } u_0 \leq \hat{u}_0,\ \phi_h(\hat{a}) \geq 0, \\ \hat{\hat{a}}, & \text{if } u_0 > \hat{u}_0,\ \phi_h(\hat{\hat{a}}) \geq 0, \\ 0, & \text{otherwise,} \end{cases}$$

where $\hat{a}$, $\hat{\hat{a}}$, $\hat{u}_0$, and $\phi_h(.)$ are as defined in Section 4.2.2. The revenue of the website in the pre-ad-block world is

$$R_{Before} = \begin{cases} N(1-\rho)[Br_b + (1-B)r_g]\hat{a}, & \text{if } u_0 \leq \hat{u}_0,\ \phi_h(\hat{a}) \geq 0, \\ N[Br_b + (1-B)r_g]\hat{\hat{a}}, & \text{if } u_0 > \hat{u}_0,\ \phi_h(\hat{\hat{a}}) \geq 0, \end{cases}$$

where ρ is as defined in Section 4.2.2.

The consumer surplus CS_{Before} in the pre-ad-block world is

$$CS_{Before} = \begin{cases} \frac{N(1-\rho)(V - V\alpha\rho + u_0)}{2} + NB\rho u_0, & \text{if } u_0 \leq \hat{u}_0,\ \phi_h(\hat{a}) \geq 0, \\ N(1-B)V - \frac{N(1-B)(V-u_0)C_r}{2C_b} + \frac{NB(V+u_0)}{2}, & \text{if } u_0 > \hat{u}_0,\ \phi_h(\hat{\hat{a}}) \geq 0, \end{cases} \tag{A.49}$$

where $\bar{\rho} = 1 - \rho$. The social welfare is equals $R_{Before} + CS_{Before}$.

Using the expressions of ad-intensities it can be easily verified that ad-block users receive an ad-light experience. Similarly, it is also straightforward to verify that the social welfare does not necessarily increase in the post-ad-block world.

Subscription Option:

Recall the setting of this extension from Section 4.6.2. As with the previous extension, it is easy to see that the revenue of the website increases in the post-ad-block world. Specifically, if the website offers to all users in the post-ad-block world the same subscription fee and ad-intensity as in the pre-ad-block world, then all the relevant equilibrium quantities – number of regular users, number of white-listers, value of the website, etc – remain the same. Thus, the revenue of the website in the pre-ad-block world is a lower bound on its revenue in the post-ad-block world.

To obtain the optimal decisions in the post-ad-block world, we partition the feasible space into two cases: (i) $v - f < u_0$ and (ii) $v - f \geq u_0$. After optimally solving each of these cases, a comparison of website's revenue yields the optimal solution for the problem.

Case (i): $v - f < u_0$. That is, $f > v - u_0$. In words, the subscription fee is high enough so that it is better for potential users to choose the outside option rather than the subscription option. Since no potential user chooses the subscription option, the analysis in this case remains the same as that without the subscription option. Thus, from our analysis in Section 4.2, the optimal decisions of the website in this case are:

$$(I_G^*,\ a_b^*,\ a_r^*) = \left(1,\ \frac{V - u_0}{C_b},\ \frac{V - u_0}{C_r}\right).$$

Also, since there is no revenue from subscription in this case, the website relies purely on advertisement revenue. Let R_{Ads}^{After} represent the website's revenue in this case. We have

$$R_{Ads}^{After} = N(V - u_0)r\left[\frac{B}{C_b} + \frac{(1 - B)}{C_r}\right]. \tag{A.50}$$

Case (ii): $v - f \geq u_0$. That is, $f \leq v - u_0$. An ad-block user chooses the subscription option if $v - f \geq v - a_b\tilde{c}_b$ or, equivalently, $\tilde{c}_b \geq f/a_b$. A regular user will choose to subscribe if $v - f \geq v - a_r\tilde{c}_r$ or, equivalently, $\tilde{c}_r \geq f/a_r$. Figure A.1 pictorially depicts these choices. The defining condition of this case (i.e., $v - f \geq u_0$) implies that all the users prefer paying for subscription over consuming the outside option. Thus, no user chooses the

Figure A.1. Ad-block users (resp., regular users) with ad-sensitivity more than f/a_b (resp., f/a_r) subscribe to the website for ad-free content.

outside option and we have $n = N$, which leads to $v = V$. Also, since all the users prefer either paying or white-listing over the outside option, it is straightforward to see that the optimal gating decision $I_G^* = 1$. We can now write the revenue of the website as

$$R_{After} = NB\left[\frac{f/a_b}{C_b} \times ra_b + \left(1 - \frac{f/a_b}{C_b}\right) f\right] + N(1-B)\left[\frac{f/a_r}{C_r} \times ra_r + \left(1 - \frac{f/a_r}{C_r}\right) f\right]$$

$$= -f^2 NB_1(a_b, a_r) + fNB_2, \tag{A.51}$$

where $B_1(a_b, a_r) = \frac{1-B}{a_r C_r} + \frac{B}{a_b C_b}$ and $B_2 = B\left(\frac{r}{C_b} + 1\right) + (1-B)\left(\frac{r}{C_r} + 1\right)$. The first-order condition gives us $f = \hat{f}(a_b, a_r) = \frac{B_2}{2B_1(a_b, a_r)}$. Using the defining condition of this case, i.e., $f \leq V - u_0$, and the concavity of R_{After} with respect to f, the optimal subscription fee $f^*(a_b, a_r) = \min\{\hat{f}(a_b, a_r),\ V - u_0\}$. Substituting this in (A.51), we have

$$R_{After}(a_b, a_r) = \begin{cases} \frac{NB_2^2}{4B_1(a_b, a_r)}, & \text{if } \hat{f}(a_b, a_r) \leq V - u_0, \\ \\ -(V-u_0)^2 NB_1(a_b, a_r) + (V-u_0)NB_2, & \text{otherwise}. \end{cases}$$

It is easy to verify that $R_{After}(a_b, a_r)$ is an increasing function of both a_b and a_r. It is reasonable to assume a practical limit, say $a_{\max}$, on each of these ad-intensities. Then, we have $a_b^* = a_r^* = a_{\max}$. Note that, in this case, the website generates revenue from both ads and subscription.

Let us now define two terms: If, in the optimal solution, a positive fraction of users choose the ad-light experience via white-listing and a positive fraction of users choose the ad-free subscription option, then we refer to such an offering by the website as a *mixed offering*. Otherwise, if no user chooses the subscription option (which, therefore, makes it redundant) then we refer to the offering by the website as an *ads-only offering*.

In Case (ii) above, it is optimal for the website to offer a mixed offering; let R_{Mix}^{After} denote the corresponding revenue. That is,

$$R_{Mix}^{After} \;=\; R_{After}(a_{\max}, a_{\max}). \tag{A.52}$$

In contrast, in Case (i) above, it is optimal for the website to choose an *ads-only offering*. The website chooses the ads-only offering if $R_{Ads}^{After} \geq R_{Mix}^{After}$; otherwise, it chooses the mixed offering. Using (A.50) and (A.52), it can be verified that both these possibilities can occur.

A similar analysis for the pre-ad-block world shows that, there too, the website can offer either the ads-only or the mixed offering. Thus, the possibilities for the website's chosen offering in the pre- and post-ad-block worlds are as follows: (A) (Ads-Only, Ads-Only), (B) (Ads-Only, Mixed), (C) (Mixed, Ads-Only), and (D) (Mixed, Mixed). When it is optimal for the website to choose the mixed offering in the post-ad-block world, then its ad-intensities for both regular and white-listers are the same (equal to $a_{\max}$). The website can thus obtain the same revenue in the pre-ad-block world too (by offering a common ad-intensity, namely $a_{\max}$, to all users). This implies that possibility (B) above, namely the website moving from an ads-only offering in the pre-ad-block world to a mixed offering in the post-ad-block world, cannot occur. Depending on the values of the parameters (e.g., $r, V,$ and u_0), each of the other three possibilities can occur. It is easy to verify that white-listers get an ad-light experience in each of these three possibilities.

Note that, in possibility (A) above, where the website continues with the ads-only offering in the post-ad-block world, our analysis of the base model applies and implies that social welfare does not necessarily increase in the post-ad-block world.

Endogenous Adoption of Ad-Blockers:

Recall the setting of this extension from Section 4.6.3. We analyze the pre- and post-ad-block worlds separately.

Post-Ad-Block Analysis: We separately analyze the two cases $I_G = 1$ and $I_G = 0$, and then compare the corresponding revenues to obtain the optimal value of I_G.

• $I_G = 1$**:** If $a_b \geq a_r$, then no user installs an ad-blocker, because there is no benefit from incurring the installation cost S. We therefore focus on the case when $a_b < a_r$. For a given set of decisions a_b and a_r, we have the following:

- ◦ A potential user installs an ad-blocker and becomes a user of the website if the following two conditions are satisfied: (i) $v - a_b\tilde{c}_a - S > v - a_r\tilde{c}_a$ or, equivalently, $\tilde{c}_a > \hat{c}$, where $\hat{c} = \frac{S}{a_r - a_b}$, and (ii) $v - a_b\tilde{c}_a - S > u_0$ or, equivalently, $\tilde{c}_a < \hat{c}_b$, where $\hat{c}_b = \frac{v - S - u_0}{a_b}$. Together, these two conditions imply that $\hat{c} < \tilde{c}_a < \hat{c}_b$.
- ◦ A potential user becomes a user of the website without installing an ad-blocker if the following conditions hold: (i) $v - a_b\tilde{c}_a - S \leq v - a_r\tilde{c}_a$, which is $\tilde{c}_a \leq \hat{c}$, and (ii) $v - a_r\tilde{c}_a \geq u_0$ or, equivalently, $\tilde{c}_a \leq \hat{c}_r$, where $\hat{c}_r = \frac{v - u_0}{a_r}$. Together, these conditions imply that $\tilde{c}_a \leq \min\{\hat{c}, \hat{c}_r\}$.

Using the above expressions for $\hat{c}$, $\hat{c}_b$, and $\hat{c}_r$, it can be easily verified that either $\hat{c} \leq \min\{\hat{c}_r, \hat{c}_b\}$ or $\hat{c} \geq \max\{\hat{c}_r, \hat{c}_b\}$. If $\hat{c} \geq \max\{\hat{c}_r, \hat{c}_b\}$, then there is no user with an ad-blocker installed. Thus, we focus on the case where $\hat{c} \leq \min\{\hat{c}_r, \hat{c}_b\}$. This condition also implies that $\hat{c}_r \leq \hat{c}_b$. Further, we can assume that the website's decisions (a_b, a_r) are such that $\hat{c}_b \leq C_a$, since the website's profit when $\hat{c}_b > C_a$ is at most that when $\hat{c}_b = C_a$. In summary, we have the following relationship among the thresholds defined above: $\hat{c} \leq \hat{c}_r \leq \hat{c}_b \leq C_a$. Potential users with ad-sensitivities in $[0, \hat{c}]$ use the website without installing an ad-blocker, potential users with ad-sensitivities in $[\hat{c}, \hat{c}_b]$ use the website with an ad-blocker installed, and potential users with ad-sensitivities in $[\hat{c}_b, C_a]$ consume the outside option. Figure A.2 depicts this segmentation of users.

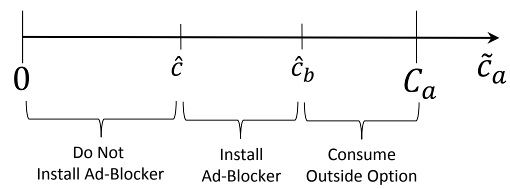

Figure A.2. Adoption of ad-blockers by users.

In the post-ad-block world, if the website gates all ad-block users and uses the same ad-intensity for white-listers as in the pre-ad-block world, then it earns the same revenue as in the pre-ad-block world. Thus, the revenue of the website in the pre-ad-block world is a lower bound on its revenue in the post-ad-block world. We now proceed to obtain the optimal ad-intensities and the social surplus.

Let n_r represent the number of regular users. Then, we have $n_r = (\hat{c}/C_a)N$. Similarly, let n_b represent the number of users who use an ad-blocker on the website. Then, we have $n_b = [(\hat{c}_b - \hat{c})/C_a]N$. Thus, the total number of users on the website can be written as $n = n_r + n_b = N(v - S - u_0)/(C_a a_b)$. Using $v = V\left(\bar{\alpha} + \alpha \frac{n}{N}\right)$, we have

$$v = \frac{V\bar{\alpha} - A_1(a_b)(S + u_0)}{1 - A_1(a_b)}, \tag{A.53}$$

where $A_1(a_b) = \frac{V\alpha}{C_a a_b}$. As in Section 4.2, we define the "value premium" for the website $\phi = v - u_0$. Recall from Section 4.2 that the website receives a positive traffic only if $\phi \geq 0$; this condition was referred to as the "survival condition". In this section, this survival condition simplifies to $V\bar{\alpha} - u_0 - S \geq 0$, and is assumed in your analysis below.

The revenue of the website can be written as

$$Revenue = \frac{Nr\big[V\bar{\alpha} - u_0 - A_1(a_b)S\big]}{C_a[1 - A_1(a_b)]},$$

from which we have

$$\frac{\partial Revenue}{\partial a_b} = -\frac{Nr(V\bar{\alpha} - u_0 - S)}{(1 - A_1(a_b))^2}\frac{V\alpha}{C_a^2 a_b^2}.$$

Since $V\bar{\alpha} - u_0 - S \geq 0$, the revenue of the website is decreasing in a_b. Thus, the optimal value of a_b is its minimum feasible value. Since $\hat{c}_b \leq C_a$ and $\hat{c}_b$ is decreasing in a_b, at the optimal value of a_b we have $\hat{c}_b = C_a$. Thus, $a_b^* = \frac{v-S-u_0}{C_a}$. Substituting a_b^* in (A.53), we have $v = V$ and, therefore, $a_b^* = \frac{V-S-u_0}{C_a}$. The condition $\hat{c} \leq \min\{\hat{c}_r, \hat{c}_b\}$ implies that $a_r \geq \frac{(V-u_0)a_b^*}{(V-u_0-S)}$. Using the derived expression for a_b^*, we have

$$a_r \geq \frac{V - u_0}{C_a}. \tag{A.54}$$

Any value of a_r that satisfies (A.54) results in an equilibrium (a_b^*, a_r); the website's revenue is the same at all these equilibria.

Since the preceding analysis corresponds to $I_G = 1$, we let R_1 refer to the optimal revenue of the website in this case. Thus, we have

$$R_1 = \frac{Nr(V - u_0)}{C_a}. \tag{A.55}$$

We now consider the other case, namely $I_G = 0$.

• $I_G = 0$**:** In this case, a user will install an ad-blocker if $v - S > v - a_r\tilde{c}_a$ or, equivalently, $\tilde{c}_a > \hat{c}_0$, where $\hat{c}_0 = \frac{S}{a_r}$. Therefore, users with an ad-sensitivity higher than $\hat{c}_0$ will not generate revenue for the website because they install ad-blockers and $I_G = 0$. Let R_0 represent the revenue of the website in this case. Then, we have

$$R_0 = \frac{\hat{c}_0}{C_a} \times N \times ra_r = \frac{NrS}{C_a}. \tag{A.56}$$

From (A.55) and (A.56), we see that $R_1 \geq R_0$. Thus, the optimal gating decision is $I_G^* = 1$.

The social surplus can now be written as

$$W_{After} = \frac{Nr(V-u_0)}{C_a} + NV - \frac{N}{2C_a}\left[a_r\hat{c}^2 + \frac{(V-S-u_0)}{C_a}(C_a-\hat{c})^2\right] - \frac{(C_a-\hat{c})}{C_a}NS. \tag{A.57}$$

We now conduct a similar analysis for the pre-ad-block world.

Pre-Ad-Block Analysis: We now conduct a similar analysis for the pre-ad-block world. In this case, the website offers a common ad-intensity, say a for all the users. A user with ad-sensitivity $\tilde{c}$ will access the website if $v - a\tilde{c} \geq u_0$ or, equivalently, $\tilde{c} \leq \frac{v-u_0}{a}$. Using this, the traffic on the website is $n = \frac{N(v-u_0)}{aC_a}$. Solving with $v = V\left(\bar{\alpha} + \alpha\frac{n}{N}\right)$, we get

$$v = \frac{V - \bar{\alpha} - u_0}{1 - D_1}, \tag{A.58}$$

where $D_1 = \frac{V\alpha}{aC_a}$. The revenue of the website in the pre-ad-block world can be written as $R_{Before} = nra = \frac{Nr(V\bar{\alpha}-u_0)}{C_a(1-D_1)}$. It is easy to see that R_{Before} is decreasing in ad-intensity a. Thus, the optimal ad-intensity a is the minimum feasible value of a. Since $\frac{v-u_0}{a} \leq C_a$, the optimal ad-intensity is $a^* = \frac{v-u_0}{C_a}$. Using (A.58), we get $v = V$ and $a^* = \frac{V-u_0}{C_a}$. The social welfare in the pre-ad-block world is

$$W_{Before} = \frac{Nr(V-u_0)}{C_a} + \frac{N(V+u_0)}{2}. \tag{A.59}$$

Using (A.57) and (A.59), it can be verified that the social surplus does not necessarily increase in the post-ad-block world. Further, a comparison of the optimal ad intensities in the pre- and post-ad-block world shows that ad-block users receive an ad-light experience.

Negative Externality (Congestion Cost) Due to Increase in Traffic:

Recall the setting of this extension from Section 4.6.4. We analyze the pre- and post-ad-block worlds separately.

Post-Ad-Block World Analysis: As with our analysis of the base model in Section 4.2, the revenue of the website in the post-ad-block world can now be written as

$$R_{After} = N(v - u_0)r\left[\frac{BI_G}{C_b} + \frac{1-B}{C_r}\right] - k_c f(n).$$

We note that, for a given set of decisions $I_G, a_b,$ and a_r, the expressions for the equilibrium traffic $n(I_G, a_b, a_r)$ and the equilibrium value $v(I_G, a_b, a_r)$ of the website remain the same as that for the base model. Thus, we have

$$v(I_G, a_b, a_r) = \frac{A_1(I_G) - u_0 A_2(a_b, a_r)}{1 - A_2(a_b, a_r)},$$

$$n(I_G, a_b, a_r) = \frac{NBI_G(v-u_0)}{C_b a_b} + (1 - I_G)NB + \frac{N(1-B)(v-u_0)}{C_r a_r}.$$

where $A_1(I_G) = V\bar{\alpha} + V\alpha(1 - I_G)B$ and $A_2(a_b, a_r) = V\alpha\left[\frac{I_G B}{C_b a_b} + \frac{1-B}{C_r a_r}\right]$.

We first comment on the comparison of the website's net revenue (i.e., the revenue minus the congestion cost) and the social surplus, and then analyze the optimal ad-sensitivities.

∘ Using the same argument as in the base model, it is easy to see that the net revenue of the website increases in the post-ad-block world. Specifically, if the website gates all ad-block users and uses the same ad-intensity in the post-ad-block world as that in the pre-ad-block world, then the website earns the same revenue and incurs the same cost. Consequently, the net revenue of the website in the pre-ad-block world is a lower bound its net revenue in the post-ad-block world.

∘ The base model is a special case, corresponding to $k_c = 0$, of the above model with congestion cost. Recall that social surplus does not necessarily increase in the post-ad-block world (Theorem 11). Therefore, the same result continues to hold.

∘ We now proceed to compare the optimal ad-intensities in the post-ad-block world. If it is optimal for the website to not gate ad-block users (i.e., set $I_G = 0$), then it follows immediately that ad-block users receive an ad-light experience (relative to the pre-ad-block world) in the post-ad-block world – this is because, in this case, ad-block users experience an ad-intensity of 0 in the post-ad-block world. Thus, we will focus on the case where $I_G = 1$.

Let $z = \frac{B}{C_b a_b} + \frac{\bar{B}}{C_r a_r}$. Using $I_G = 1$ and the expressions for v and n above, we have

$$v - u_0 = \frac{V\bar{\alpha} - u_0}{1 - V\alpha z} \text{ and} \tag{A.60}$$

$$n = \frac{N(V\bar{\alpha} - u_0)z}{1 - V\alpha z}.$$

The net revenue of the website is

$$R_{After} = \frac{N(V\bar{\alpha} - u_0)r}{1 - V\alpha z}\left[\frac{B}{C_b} + \frac{\bar{B}}{C_r}\right] - k_c\left(\frac{N(V\bar{\alpha} - u_0)z}{1 - V\alpha z}\right)^2.$$

Recall from Section 4.2 that we can restrict our search to values of (a_b, a_r) that satisfy the following constraints:

$$0 \leq v - u_0 \leq a_r C_r, \tag{A.61}$$

$$0 \leq v - u_0 \leq a_b C_b. \tag{A.62}$$

The following claim will be useful in our subsequent analysis:

Claim: *There exist optimal ad-intensities a_b^* and a_r^* are such that $a_b^* C_b = a_r^* C_r$.*

Proof of Claim: Let (a_b^*, a_r^*) denote the optimal ad-intensities. To establish the result, it suffices to construct new ad-intensities $(\hat{a}_b, \hat{a}_r)$ that also lead to the optimal net revenue. To this end, choose $\hat{a}_b$ and $\hat{a}_r$ such that

$$\hat{a}_b C_b = \hat{a}_r C_r = 1\Big/\left[\frac{B}{a_b^* C_b} + \frac{\bar{B}}{a_r^* C_r}\right]. \tag{A.63}$$

Note that $v - u_0$ and R_{After} both depend on the decisions (ad-intensities) only through z. Further, the decisions (a_b^*, a_r^*) and $(\hat{a}_b, \hat{a}_r)$ result in the same value of z. That is, $\frac{B}{a_b^* C_b} + \frac{\bar{B}}{a_r^* C_r} =$

$\frac{B}{\hat{a}_b C_b} + \frac{\bar{B}}{\hat{a}_r C_r}$. It now follows immediately that both the equilibrium value of the website and its net revenue both are the same for (a_b^*, a_r^*) and $(\hat{a}_b, \hat{a}_r)$. The result follows. ■

Using the above claim in the definition of z, we get $a_b C_b = a_r C_r = \frac{1}{z}$. This relation allows us to convert the two decision (a_b and a_r) problem into a single decision (z) problem. Using $a_b C_b = a_r C_r = \frac{1}{z}$, constraints (A.61) and (A.62) are both equivalent to $v - u_0 \leq \frac{1}{z}$. Therefore, from (A.60) we have

$$z \leq \frac{1}{V - u_0}. \tag{A.64}$$

We further simplify the expression of R_{After} be defining $\hat{z} = \frac{1}{1 - V\alpha z}$. Thus, we have $z = \left(1 - \frac{1}{\hat{z}}\right)/V\alpha$. We can now write the net revenue of the website in terms of $\hat{z}$ as

$$R_{After} = N(V\bar{\alpha} - u_0)r\hat{z}\left[\frac{B}{C_b} + \frac{\bar{B}}{C_r}\right] - k_c\left(\frac{N(V\bar{\alpha} - u_0)(\hat{z} - 1)}{V\alpha}\right)^2.$$

R_{After} is a concave function of $\hat{z}$. The first-order condition leads to $\hat{z} = \hat{z}_0$, where

$$\hat{z}_0 = \frac{rV^2\alpha^2}{2kN(V\bar{\alpha} - u_0)}\left(\frac{B}{C_b} + \frac{\bar{B}}{C_r}\right) + 1.$$

The value of z corresponding to $\hat{z} = \hat{z}_0$ is $z = z_0$, where $z_0 = \left(1 - \frac{1}{\hat{z}_0}\right)/V\alpha$. Let z^*_{post} represent the optimal value of z in the post-ad-block world. Using (A.64), we have $z^*_{post} = \min\left\{z_0, \frac{1}{V - u_0}\right\}$. Finally, using $a_b^* C_b = a_r^* C_r = 1/z^*_{post}$, we get

$$a_b^* = \frac{1}{C_b z^*_{post}}, \; a_r^* = \frac{1}{C_r z^*_{post}}.$$

Pre-Ad-Block World Analysis: As in the proof of Theorem 6, the analysis here for the pre-ad-block world can be divided in the following two cases: $0 \leq v - u_0 \leq aC_r$ and $aC_r \leq v - u_0 \leq aC_b$.

$0 \leq v - u_0 \leq aC_r$: In this case, the equilibrium expressions of v, n, and the net revenue remain the same as that in the above analysis of post-ad-block world, with an additional constraint that $a_b = a_r = a$. Thus, we now have $z = \frac{B}{C_b a} + \frac{\bar{B}}{C_r a}$. The expression of z_0

remains the same. However, in this case, z needs to satisfy the constraint $v - u_0 \le aC_r$. Substituting $v - u_0$ from (A.60) and $a = S_1/z$, where $S_1 = \frac{B}{C_b} + \frac{\bar{B}}{C_r}$, we get

$$z \le \frac{C_r S_1}{V\bar{\alpha} - u_0 + C_r S_1 V\alpha}.$$

Let z^*_{pre} represent the optimal value of z in the pre-ad-block world. Thus,

$$z^*_{pre} = \min\left\{z_0,\ \frac{C_r S_1}{V\bar{\alpha} - u_0 + C_r S_1 V\alpha}\right\}.$$

Thus, the optimal ad-intensity in the pre-ad-block world is

$$a^* = \frac{1}{z^*_{pre}}\left[\frac{B}{C_b} + \frac{\bar{B}}{C_r}\right].$$

Consider the special case where $z_0 \le \frac{1}{V - u_0}$ and $z_0 \le \frac{C_r S_1}{V\bar{\alpha} - u_0 + C_r S_1 V\alpha}$: In this case, we have $z^*_{pre} = z^*_{post} = z_0$. Thus, we have

$$a^*_b = \frac{1}{C_b z_0},\ a^*_r = \frac{1}{C_r z_0},\ \text{and } a^* = \frac{1}{z_0}\left[\frac{B}{C_b} + \frac{\bar{B}}{C_r}\right] = Ba^*_b + \bar{B}a^*_r.$$

Thus, the optimal ad-intensity in the pre-ad-block world a^* is a convex combination of the optimal post-ad-block intensities a^*_b and a^*_r. Thus, $a^* \in [a^*_b, a^*_r]$, which implies that ad-block users receive an ad-light experience in the post ad-block world.

When $z_0 > \frac{1}{V - u_0}$ or $z_0 > \frac{C_r S_1}{V\bar{\alpha} - u_0 + C_r S_1 V\alpha}$ or when $aC_r \le v - u_0 \le aC_b$, the website's optimization problem lacks sufficient structure for a clean analysis. Therefore, for simplicity, we numerically verify that ad-block users receive an ad-light experience in the post-ad-block world. Figure A.3 shows the representative behavior of the optimal ad-intensities in the pre- and post-ad-block worlds with respect to the parameter k_c of the congestion cost. Notice that the ad-intensities increase with an increase in the congestion cost. This intuition here is straightforward: as congestion cost increases, it is not necessarily desirable for the website to attract more traffic. Thus, its incentive to keep ad-intensities low (to attract more traffic) diminishes.

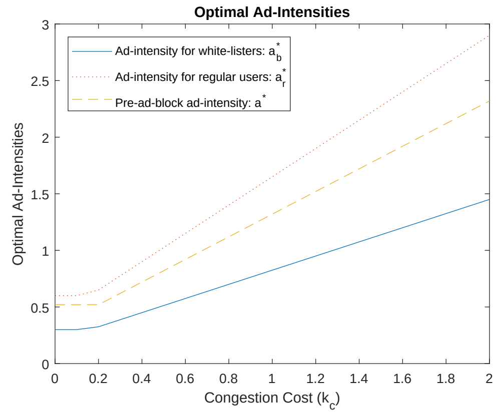

Figure A.3. Ad-intensities in the pre- and post-ad-block worlds in the presence of congestion cost: Ad-block users receive an ad-light experience in the post-ad-block world. Ad-intensities increase as the congestion-cost parameter k_c increases.

REFERENCES

Abhyankar, H. S., and S. C. Graves. 2001. Creating an Inventory Hedge for Markov-Modulated Poisson Demand: An Application and Model. *Manufacturing & Service Operations Management* 3 (4): 306–320.

Amin, K., A. Rostamizadeh., and U. Syed. 2013. Learning Prices for Repeated Auctions with Strategic Buyers. In *Proceedings of NIPS*, 1169–1177.

Andrews, M., X. Luo, Z. Fang, and A. Ghose. 2015. Mobile Ad Effectiveness: Hyper-Contextual Targeting with Crowdedness. *Marketing Science,* to appear, `http://dx.doi.org/10.1287/mksc.2015.0905`.

Arnosti, N., M. Beck, and P. Milgrom. 2016. Adverse Selection and Auction Design for Internet Display Advertising. *American Economic Review* 106 (10): 2852–2866.

Asvanund, A., K. Clay, R. Krishnan, and M. D. Smith. 2004. An Empirical Analysis of Network Externalities in Peer-To-Peer Music-Sharing Networks. *Information Systems Research* 15 (2): 155–174.

Babaioff, M., J. Hartline, and R. Kleinberg. 2009. Selling Ad Campaigns: Online Algorithms with Cancellations. In *EC 09: Proceedings of the 10th ACM Conference on Electronic Commerce*, 61–70.

Balseiro, S. R., O. Besbes, and G. Y. Weintraub. 2015. Repeated Auctions with Budgets in Ad Exchanges: Approximations and Design. *Management Science* 61 (4): 864–884.

Balseiro, S. R., J. Feldman, V. S. Mirrokni, and S. Muthukrishnan. 2014. Yield Optimization of Display Advertising with Ad Exchange. *Management Science* 60 (12): 2886–2907.

Bashyam, S., and M. C. Fu. 1998. Optimization of (s, S) Inventory Systems with Random Lead Times and a Service Level Constraint. *Management Science* 44 (12): 243–256.

Bhargava, H. K., and V. Choudhary. 2008. Research NoteWhen is Versioning Optimal for Information Goods? *Management Science* 54 (5): 1029–1035.

Bookbinder, J. H., and J.-Y. Tan. 1988. Strategies for the Probabilistic Lot-Sizing Problem with Service-Level Constraints. *Management Science* 34 (9): 1096–1108.

Central Market Research 2012. Online Advertising One of the Fastest Growing Advertising Segments. Available at: `http://finance.yahoo.com/news/Online-Advertising-One-iw-991038561.html`.

Chellappa, R. K., and S. Shivendu. 2005. Managing Piracy: Pricing and Sampling Strategies for Digital Experience Goods in Vertically Segmented Markets. *Information Systems Research* 16 (4): 400–417.

Chellappa, R. K., and S. Shivendu. 2010. Mechanism Design for "Free" but "No Free Disposal" Services: the Economics of Personalization Under Privacy Concerns. *Management Science* 56 (10): 1766–1780.

Chen, Y., W. Xue, and J. Yang. 2013. Technical Note - Optimal Inventory Policy in the Presence of a Long-Term Supplier and a Spot Market. *Operations Research* 61 (1): 88–97.

Chickering, D. M, and D. Heckerman. 2003. Targeted Advertising on the Web with Inventory Management. *Interfaces* 33 (5): 71–77.

Celis, L.E., G. Lewis, M. Mobius, and H. Nazerzadeh. 2014. Buy-It-Now or Take-a-Chance: Price Discrimination Through Randomized Auctions. *Management Science* 60 (12): 864–884.

Cidewalk, Inc. 2016. Private Communications.

Ciocan, D.F., and V. F. Farias. 2012. Dynamic Allocation Problems with Volatile Demand. *Mathematics of Operations Research* 37 (3): 501–525.

Constantin, F., J. Feldman, S. Muthukrishnan, and M. Pál. 2009. Online Ad Slotting with Cancellations. In *Proceedings of ACM-SIAM Symposium on Discrete Algorithms*, 2009, 1265–1274.

Despotakis, S., R. Ravi, and K. Srinivasan. 2017. The Beneficial Effects of Ad Blockers. *Working Paper*.

Dou, Y., M. F. Niculescu, and D. Wu. 2013. Engineering Optimal Network Effects via Social Media Features and Seeding in Markets for Digital Goods and Services. *Information Systems Research* 24 (1): 164–185.

Downey, M. 2012. Real-Time Bidding is the Next Mobile Ad Breakthrough - Here's How You Can Profit. Available at: `https://bit.ly/2GMfqkG`.

DVorkin, L. 2016. Inside Forbes: More Numbers On Our Ad Blocking Plan – and What's Coming Next. Available at: `https://www.forbes.com/sites/lewisdvorkin/2016/02/10/inside-forbes-more-numbers-on-our-ad-blocking-plan-and-whats-coming-next/#6efeea776209` (Accessed September 1, 2017).

eMarketer 2014. Total US Ad Spending to See Largest Increase Since 2004. `http://www.emarketer.com/Article/Total-US-Ad-Spending-See-Largest-Increase-Since-2004/1010982`.

Esfandiari, H., N. Korula, and V. Mirrokni. 2015. Online Allocation with Traffic Spikes: Mixing Adversarial and Stochastic Models. In *Proceedings of 16th ACM Conference on Economics and Computation*, 169–186.

Fang, Z., B. Gu, X. Luo, and Y. Xu. 2015. Contemporaneous and Delayed Sales Impact of Location-Based Mobile Promotions. *Information Systems Research*, to appear.

Feige, U., N. Immorlica, V. S. Mirrokni, and H. Nazerzadeh. 2009. "PASS Approximation," Approximation, Randomization, and Combinatorial Optimization. *LNCS 5687, Springer, 2009*, 111–124.

Feldman, J., N. Korula, V. Mirrokni, S. Muthukrishnan, and M. Pál. 2015. Online Ad Assignment with Free Disposal. In *Workshop on Internet and Network Econom. WINE '09* (Springer Verlag, Berlin), 374–385.

Forman, C., A. Ghose, and A. Goldfarb. 2009. Competition Between Local and Electronic Markets: How the Benefit of Buying Online Depends on Where You Live. *Management Science* 55 (1): 47–57.

Gallaugher, John M. 2008. Understanding Network Effects. Available at: `http://www.gallaugher.com/Network%20Effects%20Chapter.pdf` (Accessed April 6, 2018).

Garimella, K., O. Kostakis, and M. Mathioudakis. 2017. Ad-Blocking: A Study on Performance, Privacy and Counter-Measures. *arXiv preprint arXiv:1705.03193*.

Gervais, A., A. Filios, V. Lenders, and S. Capkun. 2016. Quantifying Web Adblocker Privacy. *IACR Cryptology ePrint Archive* 2016:900.

Ghosh, A., P. McAfee, K. Papineni, and S. Vassilvitskii. 2009. Bidding for Representative Allocations for Display Advertising. In *Proceedings of Workshop Internet and Network Economics*, 2009, 208–219.

Goh, K.-Y., K.-L. Hui, and I. P. Png. 2015. Privacy and Marketing Externalities: Evidence from Do Not Call. *Management Science* 61 (12): 2982–3000.

Haksöz, Ç., and S. Seshadri. 2007. Supply Chain Operations in the Presence of a Spot Market: A Review with Discussion. *Journal of the Operational Research Society* 58 (11): 1412–1429.

Hann, I.-H., K.-L. Hui, S.-Y. T. Lee, and I. P. Png. 2008. Consumer Privacy and Marketing Avoidance: A Static Model. *Management Science* 54 (6): 1094–1103.

Heilpern, W. 2016. Websites That Detect Your Ad Blocker Could Be Breaking EU Law. Available at: `http://www.businessinsider.com/ad-blocker-blockers-may-be-illegal-in-europe-2016-4` (Accessed September 1, 2017).

Jasin, S. 2014. Reoptimization and Self-Adjusting Price Control for Network Revenue Management. *Operations Research* 62 (5): 1168–1178.

Jasin, S., and S. Kumar. 2012. A Re-Solving Heuristic with Bounded Revenue Loss for Network Revenue Management with Customer Choice. *Mathematics of Operations Research* 37 (2): 313–345.

Johar, M., S. Menon, and V. Mookerjee. 2011. Analyzing Sharing in Peer-To-Peer Networks Under Various Congestion Measures. *Information Systems Research* 22 (2): 325–345.

Johnson, J. P. 2013. Targeted Advertising and Advertising Avoidance. *The RAND Journal of Economics* 44 (1): 128–144.

Kanoria, Y., and H. Nazerzadeh. 2014. Dynamic Reserve Prices for Repeated Auctions: Learning from Bids. In *Proceedings of WINE*, Vol. 232.

Katz, M. L., and C. Shapiro. 1985. Network Externalities, Competition, and Compatibility. *The American Economic Review* 75 (3): 424–440.

Kauffman, R. J., J. McAndrews, and Y.-M. Wang. 2000. Opening the "Black Box" of Network Externalities in Network Adoption. *Information Systems Research* 11 (1): 61–82.

Kesselheim, T., K. Radke, A. Tönnis, and B. Vöcking. 2014. Primal Beats Dual on Online Packing LPs in the Random-Order Model. In *Proceedings of 46th ACM Symposium on Theory of Computing*, 303–312.

Korula, N., V. Mirrokni, and H. Nazerzadeh. 2016. Optimizing Display Advertising Markets: Challenges and Directions. In *IEEE Internet Computing*, 2016.

Krammer, V. 2008. An Effective Defense Against Intrusive Web Advertising. In *Privacy, Security and Trust, 2008. PST'08. Sixth Annual Conference on*, 3–14. IEEE.

Lahiri, A., and D. Dey. 2013. Effects of Piracy on Quality of Information Goods. *Management Science* 59 (1): 245–264.

Lambrecht, A., and K. Misra. 2016. Fee or Free: When Should Firms Charge for Online Content? *Management Science* 63 (4): 1150–1165.

Lieberman, D. 2013. TV Advertising Is Surprisingly Weak Due To Internet And Economy. `http://www.deadline.com/2013/03/tv-advertising-weakness-economy-internet-analyst-report/`.

Luo, X., M. Andrews, Z. Fang, and C. W. Phang. 2013. Mobile targeting. *Management Science* 60 (7): 1738–1756.

Mahdian, M., H. Nazerzadeh and A. Saberi. 2007. Allocating Online Advertisement Space with Unreliable Estimates. In *EC '07: Proceedings of the 8th ACM Conference on Electronic Commerce*, 288–294.

Meyer, D. 2016. This Popular Ad-Blocker Keeps Beating 'Old Guard' Media in Court. Available at:
http://fortune.com/2016/03/30/adblock-plus-sueddeutsche/ (Accessed September 1, 2017).

Mirrokni, V., S. Muthukrishnan, and U. Nadav. 2010. Quasi- Proportional Mechanisms: Prior-free Revenue Maximization. In *LATIN 2010: Theoretical Informatics, LNCS 6034, Springer*, 565–576.

Nair, J., A. Wierman, and B. Zwart. 2015. Provisioning of Large-Scale Systems: The Interplay Between Network Effects and Strategic Behavior in the User Base. *Management Science* 62 (6): 1830–1841.

Najafi-Asadolahi, S., and K. Fridgeirsdottir. 2014. Cost-per-Click Pricing for Display Advertising. *Manufacturing & Service Operations Management* 16 (4): 482–497.

Niculescu, M. F., and D. J. Wu. 2014. Economics of Free Under Perpetual Licensing: Implications for the Software Industry. *Information Systems Research* 25 (1): 173–199.

Nithyanand, R., S. Khattak, M. Javed, N. Vallina-Rodriguez, M. Falahrastegar, J. E. Powles, E. De Cristofaro, H. Haddadi, and S. J. Murdoch. 2016. Ad-Blocking and Counter Blocking: A Slice of the Arms Race. In *6th USENIX Workshop on Free and Open Communications on the Internet (FOCI 16)*. USENIX Association.

Parker, G. G., and M. W. Van Alstyne. 2005. Two-Sided Network Effects: A Theory of Information Product Design. *Management Science* 51 (10): 1494–1504.

Ray, A., H. Ghasemkhani, and K. N. Kannan. 2017. Ad-Blockers: Extortionists or Digital Age Robin Hoods? *Working Paper*.

Rodriguez, A. 2017. The EU is Totally Fine with Publishers Blocking Ad Blockers. Available at: https://qz.com/882462/the-eu-is-totally-fine-with-with-publishers-blocking-ad-blockers/ (Accessed September 1, 2017).

Sheffrin, S. M. 1996. *Rational Expectations*. Cambridge University Press.

Shin, J. 2007. How Does Free Riding on Customer Service Affect Competition? *Marketing Science* 26 (4): 488–503.

Shriver, S. K., H. S. Nair, and R. Hofstetter. 2013. Social Ties and User-Generated Content: Evidence from an Online Social Network. *Management Science* 59 (6): 1425–1443.

Song, J.-S., and P. Zipkin. 1993. Inventory Control in a Fluctuating Demand Environment *Operations Research* 41 (2): 351–370.

Storey, G., D. Reisman, J. Mayer, and A. Narayanan. 2017. The Future of Ad Blocking: An Analytical Framework and New Techniques. *arXiv preprint arXiv:1705.08568*.

Stühmeier, T., and T. Wenzel. 2011. Getting Beer During Commercials: Adverse Effects of Ad-Avoidance. *Information Economics and Policy* 23 (1): 98–106.

Tayur, S., R. Ganeshan, and M. Magazine. 2012. *Quantitative Models for Supply Chain Management*, Volume 17. Springer Science & Business Media.

Turner, J. 2014. The Planning of Guaranteed Targeted Display Advertising. *Operations Research* 60 (1): 18–33.

Vallade, J. 2008. Adblock Plus and the Legal Implications of Online Commercial-Skipping. *Rutgers L. Rev.* 61:823.

Varian, H. R. 1997. Versioning Information Goods.

Wall Street Journal 2014. What Do Consumers Hate More Than TV Ads? Online Video Commercials. Available at: `http://blogs.wsj.com/cmo/2014/07/15/what-do-consumers-hate-more-than-tv-ads-online-video-commercials/` (Accessed September 1, 2017).

Wall Street Journal 2015. Apples Ad Blockers Rile Publishers. Available at: `http://www.wsj.com/articles/apples-ad-blockers-raise-tensions-1440974849` (Accessed September 1, 2017).

Wall Street Journal 2017. Google Plans Ad-Blocking Feature in Popular Chrome Browser. Available at: `https://www.wsj.com/articles/google-plans-ad-blocking-feature-in-popular-chrome-browser-1492643233` (Accessed September 1, 2017).

Wang, Y., and Y. Gerchak. 1996. Periodic Review Production Models with Variable Capacity, Random Yield, and Uncertain Demand. *Management Science* 42 (1): 130–137.

Wei, X. D., and B. R. Nault. 2014. Monopoly Versioning of Information Goods When Consumers Have Group Tastes. *Production and Operations Management* 23 (6): 1067–1081.

Wu, S.-y., and P.-y. Chen. 2008. Versioning and Piracy Control for Digital Information Goods. *Operations Research* 56 (1): 157–172.

Yano, C. A., and H. L. Lee. 1995. Lot Sizing with Random Yields: A Review. *Operations Research* 43 (2): 311–334.

Zhang, K., T. Evgeniou, V. Padmanabhan, and E. Richard. 2012. Content Contributor Management and Network Effects in a UGC Environment. *Marketing Science* 31 (3): 433–447.

Zhang, W., S. Yuan, and J. Wang. 2014. Optimal Real-Time Bidding for Display Advertising. *Proceedings of 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1077–1086.

BIOGRAPHICAL SKETCH

Manmohan was born in Jodhpur, India. Before joining the PhD program at UT-Dallas, he completed B.Tech in Electrical Engineering from Indian Institute of Technology Kanpur. He is interested in teaching Business Intelligence and Analytics, IT Services Management, Enterprise Systems, Real-Time Digital Advertising, Database Management. His primary research interests fall under the broad umbrella of the digital advertising ecosystem. Within this context, he has investigated several challenges faced by different stakeholders, including (i) bidding strategies for mobile-promotion platforms in real-time ad-auctions, (ii) decisions of publishers on whether to allow ad-free access to ad-block users or require them to white-list, and the corresponding ad-intensities, and (iii) experimentation strategies for advertisers in real-time advertising. He also has a strong research interest in analyzing the impact of disruptive technological innovations on traditional players. For instance, he is investigating the impact of Massive Open Online Courses (MOOCs) on the traditional players in the education industry.

CURRICULUM VITAE

Manmohan Aseri

2200 Waterview Parkway, Apt 1738,
Richardson, TX 75080

Mobile: (469) 360-2466
E-mail: manmohan.aseri@utdallas.edu

Research Interests

Real-Time Digital Advertising, Ad-Auctions, Experimentation Strategies, Disruptive Innovations.

Research Methodologies

Game Theory, Optimization, Econometrics.

Teaching Interests

Business Intelligence and Analytics, IT Services Management, Enterprise Systems, Real-Time Digital Advertising, Database Management.

Education

2012 – present	**The University of Texas at Dallas, USA** Naveen Jindal School of Management Ph. D. Candidate in Information Systems **Advisor: Vijay Mookerjee**
2003 – 2007	**Indian Institute of Technology Kanpur, India** B. Tech. in Electrical Engineering

Research Papers

- Aseri, Manmohan, M. Dawande, G. Janakiraman, and V. Mookerjee. "Procurement Policies for Mobile-Promotion Platforms." **Forthcoming in *Management Science***.
  - **Best Paper Runner-Up Award in CIST 2016.**
- Aseri, Manmohan, M. Dawande, G. Janakiraman, and V. Mookerjee. "Ad-Blockers: A Blessing or a Curse?" **Major Revision in *Information Systems Research***.
  - **Best Conference Paper Winner in CIST 2017.**
  - **Nominated for Best Student Conference Paper Award in CIST 2017.**
- Aseri, Manmohan, M. Dawande, G. Janakiraman, and V. Mookerjee. "Experimentation in Ad-Campaigns." ***Work in Progress***.
- Aseri, Manmohan, H. Singh, and V. Mookerjee. "Impact of Open Online Courseware on Learners and Providers: An Empirical Study." ***Work in Progress***.

Teaching Experience at UT Dallas

Spring '17	*Instructor* for Introduction to Programming in Java (55 Students) ◦ Evaluation: 4.91/5.
Fall '15	*Instructor* for Information Technology for Business (48 Students) ◦ Evaluation: 4.87/5.
Summer '17	*Instructor* for Information Technology for Business (46 Students) ◦ Evaluation: 4.64/5.
Fall '16	*Instructor* for Introduction to Programming in Java (48 Students) ◦ Evaluation: 4.51/5.

Teaching Assistant

Information Technology Security, Database Management, Information Technology for Business, Introduction to Programming in Java.

Work Experience

Sep 2007 – May 2008	**ARM Embedded Technologies Pvt. Ltd. Bangalore, India.** **Position:** Software Engineer at ARM **Project:** Development of debugger for ARM platforms. **Tools:** (i) ARM assembly language and (ii) Perl.
May 2008 – Dec 2008	**NXP Semiconductors. Bangalore, India.** **Position:** Software Engineer at NXP's Telecom Division. **Project:** Development of UMTS protocol stack. **Tools:** (i) C language on Unix platform, (ii) Perl and Shell scripting for automation, and (iii) ClearCase for version control.
Feb 2009 – May 2011	**Samsung Electronics Pvt. Ltd. Noida, India.** **Position:** Software Engineer at Samsung's Mobile Division. **Project:** Implementation of SyncML protocol for data synchronization of mobiles with servers. **Tools:** (i) C language on Windows platform, (ii) Perforce for version control, and (iii) Visual Studio for debugging.
May 2011 – May 2012	**Indian Institute of Technology, Delhi.** **Position:** Research Associate. **Project:** Allocation of channels in Cognitive Radio Networks. **Tool:** C language on Windows platform.

Conference Presentations

- ICIS, Fort Worth, December, 2015 (Poster Presentation).
- POMS Annual Conference, Orlando, May, 2016.
- CIST, Nashville, November, 2016.
- INFORMS Annual Meeting, Nashville, November, 2016.
- POMS Annual Conference, Seattle, May, 2017.
- CIST, Houston, October, 2017.
- INFORMS Annual Meeting, Houston, October, 2017 (two presentations).

Invited Talks

- Tepper School of Business at Carnegie Mellon University.
- Krannert School of Management at Purdue University.
- Muma College of Business at University of South Florida.

Academic Services and Awards

- Reviewer for *Information Systems Research (ISR), Production and Operations Management (POM) Decision Support Systems (DSS), International Conference on Web Services (ICWS)*.
- **Winner:** JSOM Three-Minute Dissertation Competition.
- Organized a session on digital advertising in INFORMS annual meeting 2017.

Selected Coursework

Econometrics	Statistics
Advanced Managerial Economics	Game Theory
Optimization	Optimal Control
Advanced Topics in Knowledge Management	Probability and Stochastic Processes
PhD Seminars on Empirical Research	PhD Seminars on Game Theory

Computer Skills

- Data Analysis Tools: R, SAS, Matlab.
- Programming Languages: C, Java, Python, Perl, SQL.

References

Vijay Mookerjee
Charles and Nancy Davidson Professor of Information Systems and Operations Management
Naveen Jindal School of Management
The University of Texas at Dallas, Richardson TX
✉ vijaym@utdallas.edu
Office Phone: (972) 883-4414

MILIND DAWANDE
Ashbel Smith Professor of Operations Management
Naveen Jindal School of Management
The University of Texas at Dallas, Richardson TX
✉ milind@utdallas.edu
Office Phone: (972) 883-2793

GANESH JANAKIRAMAN
Ashbel Smith Professor of Operations Management
Naveen Jindal School of Management
The University of Texas at Dallas, Richardson TX
✉ ganesh@utdallas.edu
Office Phone: (972) 883-5846

HARPREET SINGH
Associate Professor of Information Systems
Naveen Jindal School of Management
The University of Texas at Dallas, Richardson TX
✉ harpreet@utdallas.edu
Office Phone: (972) 883-4770