Essays on Optimization and Incentive Contracts

by

Pranava Raja Goundan

Submitted to the Sloan School of Management
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Operations Research

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2007

Author . . . .
Sloan School of Management
May 17, 2007

Certified by . . . .
Andreas S. Schulz
Associate Professor of Operations Research
Thesis Supervisor

Certified by . . . .
David Simchi-Levi
Professor of Engineering Systems
Thesis Supervisor

Accepted by . . . .
Dimitris Bertsimas
Boeing Professor of Operations Research
Codirector, Operations Research Center
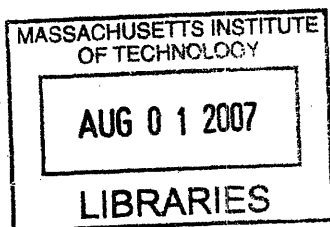
Essays on Optimization and Incentive Contracts

by

Pranava Raja Goundan

Submitted to the Sloan School of Management
on May 17, 2007, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Operations Research

Abstract

In this thesis, we study important facets of two problems in methodological and applied operations research. In the first part of the thesis, motivated by optimization problems that arise in the context of Internet advertising, we explore the performance of the greedy algorithm in solving submodular set function maximization problems over various constraint structures. Most classic results about the greedy algorithm assume the existence of an optimal polynomial-time incremental oracle that identifies in any iteration, an element of maximum incremental value to the solution at hand. In the presence of only an approximate incremental oracle, we generalize the performance bounds of the greedy algorithm in maximizing nondecreasing submodular functions over special classes of matroids and independence systems. Subsequently, we unify and improve on various results in the literature for problems that are specific instances of maximizing nondecreasing submodular functions in the presence of an approximate incremental oracle. We also propose a randomized algorithm that improves upon the previous best-known 2-approximation result for the problem of maximizing a submodular function over a partition matroid.

In the second part of the thesis, we focus on the design and analysis of simple, possibly non-coordinating contracts in a single-supplier, multi-retailer supply chain where retailers make both pricing and inventory decisions. Specifically, we introduce a buy-back menu contract to improve supply chain efficiency, and compare two systems, one in which the retailers compete against each other, and another in which the retailers coordinate their decisions to maximize total expected retailer profit. In a linear additive demand setting, we show that for either retailer configuration, the proposed buy-back menu guarantees the supplier, and hence the supply chain, at least 50% of the optimal global supply chain profit. In particular, in a coordinated retailers system, the contract guarantees the supply chain at least 75% of the optimal global supply chain profit. We also analyze the impact of retail price caps on supply chain performance in this setting.

Thesis Supervisor: Andreas S. Schulz
Title: Associate Professor of Operations Research

Thesis Supervisor: David Simchi-Levi
Title: Professor of Engineering Systems

Acknowledgments

The journey to this thesis represents an important chapter of my life, and in many ways, the pinnacle of my academic education so far. I therefore first express my sincerest gratitude to all my teachers – past, present, and future, for all that they have sought to teach me. Any lessons that remain to be learnt reflect my own shortcomings.

I am deeply indebted to my advisors, Prof. Andreas Schulz and Prof. David Simchi-Levi, for their guidance, support, and encouragement. Over the past five years, Andreas has been a virtuous mentor, unwavering supporter and a good friend, under whose tutelage I have learnt to set my own course of research and in the end, have found problems that I have come to love. I thank him for sharing with me his insights and humor over many research sessions and badminton games. I began working with David three years ago, and his enthusiasm for that new result has never ceased to motivate and amaze me ever since. His wisdom and advice have proved invaluable over the course of my doctoral studies and for that, I am grateful. It has truly been an honor and a pleasure to work with them both.

Special thanks are also due to Prof. Georgia Perakis, under whose direction I gained my first teaching experience at MIT. Ever since then, Georgia has been very supportive of my academic progress, and her counsel and insights have helped greatly improve my research. Georgia was also kind enough to be a part of my thesis committee.

I would also like to thank Prof. Lap Mui Ann Chan and Prof. Jim Orlin, with whom I have had the opportunity to collaborate during my stay at MIT. I have benefited greatly from my interactions with them and remain appreciative of them for their support.

I remain forever grateful to my undergraduate mentor, Prof. T. T. Narendran, for introducing me to the world of operations research and encouraging me to pursue graduate studies at the Operations Research Center at MIT.

The ORC at MIT is a most wondrous world in itself. What makes the ORC such a special place is the students, staff and faculty that are a part of it. Paulette, Laura, Andrew, and Veronica (before Andrew) keep this world in order while the faculty and students go about their "revolutions". My sincerest thanks to them for their efforts.

Contents

List of Figures

Chapter 0

Outline of Thesis

This thesis concerns itself with two diverse topics that arise in optimization and incentive contracts, respectively: *submodular set function maximization*, and *incentive contracts in a pricing regime*. For this reason, the thesis is structured in two parts, PART I and PART II. In this chapter, we briefly outline the contents of the thesis.

Part I

In **Chapter 1**, we consider optimization problems that arise in the context of Internet advertising. There are two predominant paradigms in web advertising: *display-based advertising* and *search-based advertising*. Display-based advertising broadly refers to the paradigm of displaying ads on webpage banners, rich media ads such as videos, and graphics that pop up on webpages. Search-based advertising refers to the format of advertising based on sponsored search results for search queries that users type in search engines. We analyze a few problems that have been studied in the literature, covering both advertising paradigms. Interestingly, we observe that a greedy algorithm is an efficient heuristic in solving many variants of advertisement placement problems. Based on this observation, we find that such problems may actually be posed in the more general framework of maximizing a nondecreasing submodular set function over independence system constraints. This would provide the basis therefore, for the classical results of Fisher, Nemhauser, Wolsey, and others [19, 38, 80], on the performance of the greedy algorithm for submodular function maximization, to apply. However,

13

these results about the greedy algorithm assume the existence of an optimal polynomial-time incremental oracle that identifies in any iteration, an element of maximum incremental value to the solution at hand. In the problems that we consider, it turns out that selecting an element of maximum incremental value may itself be a hard problem, such as the KNAP-SACK problem or the RECTANGLE PACKING problem. Therefore, this raises the question of whether the aforementioned results for the greedy algorithm would generalize when only an *approximate* incremental oracle is available for the greedy algorithm. This is precisely the question we address in Chapter 2.

In **Chapter 2**, given only an approximate incremental oracle, we generalize the performance bounds of the greedy algorithm, and a related variant, the *locally greedy algorithm* [38], in maximizing nondecreasing submodular functions over (i) uniform matroids, (ii) partition matroids, and (iii) independence systems. Subsequently, we are able to unify and reinterpret several results in the recent literature by showing that such problems are in fact special cases of maximizing nondecreasing submodular functions in the presence of an approximate incremental oracle. In the case of certain problems, we are even able to improve on the best-known approximation results for these problems.

Based on these insights, we also develop an improved randomized $(2 - \frac{1}{n})$-approximation algorithm for the problem of maximizing a nondecreasing submodular function over a partition matroid, where n is the maximum number of elements in any partition. Since MAXimum SATisfiability problem is an example of maximizing a submodular function over a partition matroid with $n = 2$, this algorithm reveals, to the best of our knowledge, a new $\frac{3}{2}$-approximation algorithm for it. We also reinterpret the greedy algorithm as a limiting case of this randomized algorithm.

Chapters 1 and 2 are based on joint work with Prof. Andreas S. Schulz, and reflect the chronological order in which this work was done. For the ease of the reader, the chapters have also been kept self-contained and may be read in a stand-alone manner.

Part II

A supply chain setting is inherently characterized by strategic interactions between multi-

ple agents seeking to maximize their own utilities. The design of coordinating contractual agreements as incentives to align the interests of different members of a supply chain has received wide attention in the literature. In **Chapter 3**, we focus on the design and analysis of simple, possibly *non-coordinating* contracts in a single-supplier, multi-retailer supply chain where retailers make both pricing and inventory decisions. Specifically, we introduce a buy-back type incentive mechanism, known as a *buy-back menu contract*, to improve supply chain efficiency. We compare two systems, one in which the retailers compete against each other and another in which the retailers coordinate their decisions to maximize total expected retailer profit. In a linear additive demand setting, we show that for either retailer configuration, the proposed buy-back menu guarantees the supplier, and hence the supply chain, at least 50% of the optimal global supply chain profit. In particular, in a coordinated retailers system, the contract guarantees the supply chain at least 75% of the optimal global supply chain profit. We also analyze the impact of retail price caps on supply chain performance in this setting, and establish that price caps can hurt supply chain performance, while being significantly detrimental to retailers.

Chapter 3 is based on joint work with Prof. Lap Mui Ann Chan and Prof. David Simchi-Levi.

PART I

Chapter 1

Optimization in Internet Advertising

1.1. Introduction

The advent of the Internet as a global communications network for a variety of services has made it an attractive medium for the publicity and advertising of commercial products and services. Advertising on the web has grown tremendously from a $ 906.5 million per annum market in 1997 to a $ 12.5 billion per annum market in 2005 (Interactive Advertising Bureau [11]). Industry research firms estimate that continued reallocation of traditional media advertising spending to the Internet will see web advertising spending grow to over $16 billion in 2006 (Wall Street Journal [1]). Many firms, such as Google, Yahoo, and MSN, to name a few, have for this very reason adopted a business model where revenue is generated from selling advertisement slots while they provide free content-based service to their clients. Hence, maximization of revenue from advertisements is of utmost interest to such firms. For this reason, optimization in advertising has a been a problem of interest in recent times [4, 5, 28, 37, 55].

Web advertisements are used for two primary purposes: *direct marketing* and *brand awareness*. According to a study by Nielsen/NetRatings [81], web advertisers use a wide selection of ad formats to convey their messages. However, two predominant formats of web advertising, namely *search-based advertising* and *display-based advertising*, account for over

65% of the revenue generated (Interactive Advertising Bureau [11]). In the current work, we focus on optimization problems that arise in the context of these two advertising formats.

Display-based advertising broadly refers to the paradigm of displaying ads on webpage banners, rich media ads such as videos, and graphics that pop up on webpages. Despite a growing trend towards a search-based advertising paradigm, display-based banner advertisements remain an important form of online advertising, accounting for over 21% of all web advertising revenue in 2005 (Interactive Advertising Bureau [11]). One of the distinctive features of banner advertising is the *space-sharing* of these advertisement banner spaces, wherein possibly two or more advertisements of varying sizes share the banner space during a particular time slot. Also, in order to increase advertising effectiveness, the banner space is *time-shared*. In other words, the same banner space is updated periodically with different ads.

One of the most popular pricing models for web advertisements is the cost-per-thousand-impressions (CPM) model (see [52, 55, 11]), associated with over 46% of web advertising revenue in 2005 (Interactive Advertising Bureau [11]). In this model, the cost of advertising is associated with the number of exposures of the advertisement. This model may also be improved by providing guarantees on the number of exposures in a given period of time. In this paper, we study the advertisement revenue maximization problem under such a pricing scheme. In subsequent sections, we discuss a few variants of a problem wherein a website owner would like to decide which advertisements it would like to display on a banner space of fixed dimensions, over a period of time. We broadly refer to this problem and its variants as the AD PLACEMENT problem.

Characteristic of the fast-evolving nature of the Internet, paradigms of web advertising have also shifted rapidly in the past few years. Search-based advertising refers to the format of advertising based on sponsored search results for search queries that users type in search portals. Search-based advertising is increasingly the most important channel of Internet advertising, generating 41% of all web advertising revenue in 2005 [11], and is the main revenue stream for Internet search companies such as Google and Yahoo. Moreover, pricing formats

for advertising having also recently been trending towards *performance-based schemes* such as cost-per-click (CPC) and cost-per-action (CPA) [37, 11, 49]. Later in this chapter, we also discuss briefly some optimization problems that arise within this setting, and show that even though paradigms may have shifted, the underlying structure of web advertising optimization problems remain closely connected to submodular function maximization problems, which is our focus in Chapter 2.

1.2. Preliminaries

Before we describe the optimization problems under consideration, we briefly review some concepts and optimization problems of relevance to our study. We begin with the notion of an approximation algorithm. An α-**approximation algorithm** for a maximization problem $\mathcal{P}$ is a polynomial-time algorithm, A, that for all instances I of problem, $\mathcal{P}$, generates a feasible solution of value $A(I)$ such that $OPT(I) \leq \alpha \cdot A(I)$, where $OPT(I)$ is the objective value of an optimal solution to instance I. Observe that by this definition, it must be that $\alpha \geq 1$. A **fully polynomial-time approximation scheme** (FPTAS) provides, for every $\epsilon > 0$, a $(1+\epsilon)$-approximation algorithm whose running time is polynomial in both the size of the input and $1/\epsilon$. More generally, a **polynomial-time approximation scheme** (PTAS) provides a $(1 + \epsilon)$ approximation algorithm whose running time is polynomial in the size of the input, for any constant ϵ.

1.2.1 Knapsack Problem

KNAPSACK

Instance: Nonnegative integers n, $p_1, \ldots, p_n$, $w_1, \ldots, w_n$ and W.

Task: Find a subset $S \subseteq \{1, \ldots, n\}$ such that $\sum_{j \in S} w_j \leq W$ and $\sum_{j \in S} p_j$ is maximum.

While the KNAPSACK problem is known to be NP-hard, it belongs, in some sense, to the "easiest" class of NP-hard problems. It admits a pseudopolynomial time dynamic-programming based algorithm. Ibarra and Kim [56] extended the idea of this DP-based algorithm to de-

19

velop a FPTAS for the KNAPSACK problem, whose running time is $O(\frac{n^3}{\epsilon})$. This is the best possible algorithmic result one would hope to expect for this problem, unless P=NP. Since then, more efficient FPTASes have been developed for the KNAPSACK problem, but for the sake of simplicity, we will use the above mentioned FPTAS in our analysis. We refer the reader to Martello and Toth [73] or Kellerer et al. [62] for a detailed survey on this problem and its variants.

1.2.2 Multiple Knapsack Problem

MULTIPLE KNAPSACK

Instance: Nonnegative integers, n, m, $p_1, \ldots, p_n$, $w_1, \ldots, w_n$, and $W_1, \ldots, W_m$.

Task: Find m subsets $S_1, \ldots, S_m \subseteq \{1, \ldots, n\}$, $S_i \cap S_k = \emptyset$ for $i \neq k$, such that $\sum_{j \in S_i} w_j \leq W_i$ for $i = 1, \ldots, m$ and $\sum_{i=1}^{m} \sum_{j \in S_i} p_j$ is maximum.

The MULTIPLE KNAPSACK problem (MKP) is an extension of the classical knapsack problem. It is known to be strongly NP-hard, thus precluding the existence of an FPTAS for it. For the case of the MKP with identical bin capacities, a PTAS was constructed by Kellerer [61]. Subsequently, the above problem was also shown to have a PTAS by Chekuri and Khanna [15]. They also demonstrated that the above problem is a restricted version of the GENERALIZED ASSIGNMENT PROBLEM (GAP), for which a 2-approximation algorithm was proposed by Shmoys and Tardos [94]. However, Chekuri and Khanna [15] also showed that even certain restricted cases of GAP, more general than MKP, are APX-hard and hence have no chance of having a PTAS unless P=NP. The reader is referred to Kellerer et al. [62] for a detailed survey of this problem.

1.2.3 Rectangle Packing Problem

The RECTANGLE PACKING problem is an extension of the KNAPSACK problem to two dimensions. Informally stated, the objective is to select a subset of rectangles of maximum weight that may be *feasibly* packed in a larger rectangle. In a feasible packing of rectangles,

as termed above, none of the rectangles may overlap. Also, none of the rectangles may be rotated in a packing. More formally, the problem is stated as follows:

RECTANGLE PACKING

Instance: Set of n rectangles, $R_i = (a_i, b_i, p_i)$, where $a_i \leq a$ and $b_i \leq b$ are the height and width of R_i respectively and p_i is the profit associated with R_i. Also, a big rectangle R of height a and width b.

Task: Find a subset of rectangles $S \subseteq \{R_1, \ldots, R_n\}$ that can be feasibly packed in R which maximizes $\sum_{i \in S} p_i$.

The RECTANGLE PACKING problem is known to be strongly NP-hard even for packing squares with identical profits. The best-known approximation result for the RECTANGLE PACKING problem is due to Jansen and Zhang [58], who develop a $(2 + \epsilon)$-approximation algorithm for this problem based on a strip-packing result of Steinberg [97].

1.2.4 Submodular Functions and Matroids

A real-valued set function $f : 2^E \to \mathbb{R}$ is **normalized, nondecreasing** and **submodular** if it satisfies the following conditions, respectively:

(F0) $f(\emptyset) = 0$;

(F1) $f(A) \leq f(B)$ whenever $A \subseteq B \subseteq E$;

(F2) $f(A) + f(B) \geq f(A \cup B) + f(A \cap B)$ for all $A, B \subseteq E$, or equivalently:

(F2a) $f(A \cup \{e\}) - f(A) \geq f(B \cup \{e\}) - f(B)$ for all $A \subseteq B \subseteq E$ and $e \in E \setminus B$, or equivalently:

(F2b) $f(A \cup C) - f(A) \geq f(B \cup C) - f(B)$ for all $A \subseteq B \subseteq E$ and $C \subseteq E \setminus B$.

Henceforth, whenever we refer to submodular functions, we shall in particular imply normalized, nondecreasing, submodular functions. We also adopt the following notation: For any two sets $A, B \subseteq E$, we define the **marginal value** (incremental value) of set A to set B as

$$\rho_A(B) = f(A \cup B) - f(B).$$

Additionally, we will use the subscript e instead of $\{e\}$ whenever the context is clear. In particular, **(F2a)** can equivalently be written as $\rho_e(A) \geq \rho_e(B)$ for $A \subseteq B$.

A set system $(E, \mathcal{F})$, where E is a finite set and $\mathcal{F}$ is a collection of subsets of E, is an **independence system** if it satisfies the following properties:

(M1) $\emptyset \in \mathcal{F}$;

(M2) If $X \subseteq Y \in \mathcal{F}$ then $X \in \mathcal{F}$.

An independence system $(E, \mathcal{F})$ is a **matroid** if it satisfies the additional property:

(M3) If $X, Y \in \mathcal{F}$ and $|X| > |Y|$, then there is an $x \in X \setminus Y$ with $Y \cup \{x\} \in \mathcal{F}$.

In this chapter, we will also focus our attention on a few special types of matroids. A **uniform matroid** is specified by $(E, \mathcal{F})$, where E is a finite set, k is a positive integer, and

$$\mathcal{F} := \{F \subseteq E : |F| \leq k\}.$$

A **partition matroid** is specified by $(E, \mathcal{F})$, where $E = \cup_{i=1}^{k} E_i$ is the disjoint union of k sets, $l_1, \ldots, l_k$ are positive integers, and

$$\mathcal{F} = \{F : F = \cup_{i=1}^{k} F_i \text{ where } F_i \subseteq E_i, |F_i| \leq l_i \text{ for } i = 1, \ldots, k\}.$$

We refer the reader to standard combinatorial optimization textbooks [21, 42, 67, 79, 93] for a substantive discussion on these topics.

1.3. Our Results

In this chapter, we study variants of optimization problems that arise in Internet advertising. Our focus is primarily directed to problems related to *display-based* banner advertisements, namely AD PLACEMENT and related problems. In Section 1.4, we describe the AD PLACEMENT problem more formally. We then survey related work done on the problem and past results in Section 1.5. In Section 1.6, we formulate our problem as an integer program, and present a proof for its strong NP-hardness. With an aim to develop improved approximation

results, in Section 1.7, we study a relaxation of the AD PLACEMENT problem and show that the relaxed problem has a relevant interpretation in the context of banner advertisements as well. In Section 1.7.1, we propose a greedy algorithm, and a polynomial-time variant of it, to solve the relaxed AD PLACEMENT problem. In Section 1.7.2, we show that this algorithm is in fact a $(\frac{e}{e-1} + \epsilon) \approx (1.58 + \epsilon)$-approximation algorithm for the relaxed problem, which as we show in Chapter 2, leads us to also prove some interesting generalizations of classical results on the performance of the greedy algorithm to maximize submodular functions over matroids and independence systems. These results are then extended for a two-dimensional variant of the AD PLACEMENT problem in Section 1.8, and we present an $(\frac{\sqrt{e}}{\sqrt{e}-1} + \epsilon) \approx (2.54 + \epsilon)$-approximation bound for the relaxed 2-DIMENSIONAL AD PLACEMENT PROBLEM.

We briefly discuss a few optimization problems that arise in the context of *search-based advertising* in Section 1.9. Specifically, from an search portal's perspective, based on the AdWords paradigm of Google, we consider the ADWORDS ASSIGNMENT proposed by Fleischer et al. [39]. We show that this problem is also closely related to the submodular function maximization problems that AD PLACEMENT was a special case of. Using this insight, we later improve on a $(3.16 + \epsilon)$-approximation result for ADWORDS ASSIGNMENT [39] and construct a greedy $(3 + \epsilon)$-approximate algorithm for the problem. We also discuss an advertiser's BUDGET OPTIMIZATION problem in an AdWords auction setting, that was recently studied by Feldman et al. [37]. We finally conclude in Section 1.10 with some closing remarks.

1.4. Ad Placement: Problem Description

The AD PLACEMENT problem was first studied by Adler et al. [4], and was referred to by them as the offline version of the Ad Placement problem. The problem may be defined more formally as follows: We are given a rectangular display area of fixed height (for eg., a banner on a web page) and width W which needs to be utilized for T periods, which we model as T bins of size W each. Also, we are given a collection of N prospective ads, all of which have the same height, which is the height of the display area, but may have differing widths. These N ads are modeled as jobs, each defined by a triple (c_i, t_i, w_i). Here, t_i represents the

number of times that ad i must appear in T periods (*exposure*), w_i the size of the i^{th} ad and c_i represents the profit obtained for a *feasible assignment* of all t_i copies of the i^{th} ad. A feasible assignment of ads consists of k ($0 \leq k \leq N$) sets of jobs, all of whose copies have been assigned to the T bins, subject to the constraints that:

(a) no two copies of the same job are assigned to the same bin, and

(b) the capacity of all jobs assigned to a bin does not exceed the capacity of the bin itself.

The advertising scheduler is to decide which of these ads to accept or reject, and if accepted, how to place these ads in a feasible manner on the T banners, in order to maximize the total profit obtained. More formally, the problem may be stated as follows:

AD PLACEMENT

Instance: Ads $A_1, A_2, \ldots, A_N$, each A_i specified by (c_i, t_i, w_i), for $i = 1, \ldots, N$; a banner of size W ; T time slots. ($T \geq t_i$, for $i = 1, \ldots, N$)

Task: Find a feasible allocation of ads with maximum profit.

We observe that the notion of *time-sharing* is captured by the exposure specification of each ad, that we model as copies of the ads to be placed among T different time slots. Also, the differing ad sizes, given by w_i for each ad, means that two or more different ads may be placed simultaneously on the banner, as long as their combined width does not exceed the width of the banner itself. Thus we utilize the concept of *space-sharing* of the banner space.

In stating the AD PLACEMENT problem, we have made the following assumptions:

- The advertiser is forbidden to place more than one copy of the same ad in any banner at the same time, as this inherently decreases the exposure of the commodity being advertised. In fact, it is this assumption that differentiates this problem from related scheduling problems and the MULTIPLE KNAPSACK problem.

- The exposure of every ad may be specified exactly in the form $\frac{t_i}{T}$, with t_i being an integer (clearly, $t_i \leq T$, for $i = 1, \ldots, N$). In other words, T may be viewed as the least common multiple of the denominators of the exposure fraction of all the ads.

- We cannot obtain partial profits by the partial fulfillment of ad requirements, i.e., by scheduling fewer than t_i copies of ad i. However, in Section 1.7, we relax this assumption to study an alternate variant of the AD PLACEMENT problem.

- Finally, the profit obtained from scheduling an ad need not depend on only the size of the ad, and might also depend on factors such as product competition, desire for visibility on the particular banner/website, etc. This general assumption was first studied in the context of the AD PLACEMENT problem by Freund and Naor [40]. We note that this assumption is more general than those made by Adler et al. [4], Kumar et al. [28], Amiri and Menon [7] and Dean and Goemans [30], who assumed that $c_i = w_i$.
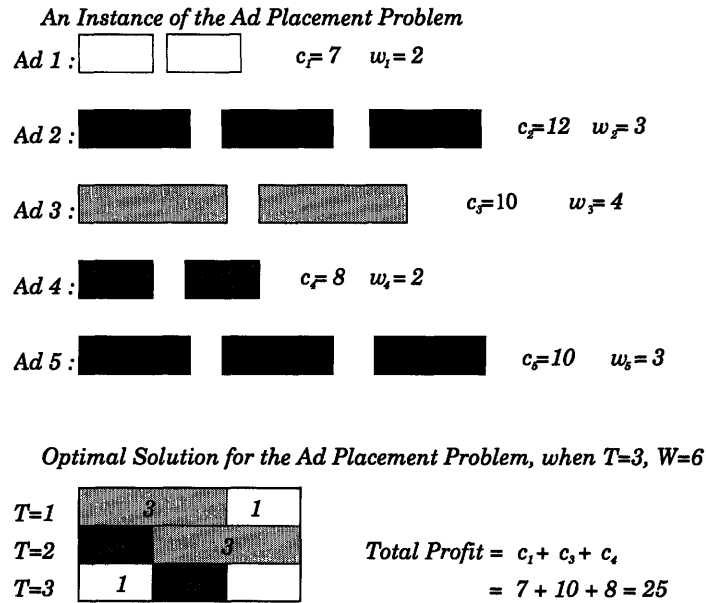
Figure 1-1: An instance of the AD PLACEMENT problem

We survey these and other results related to the AD PLACEMENT problem in Section 1.5. Note that if in the AD PLACEMENT problem, we assume that $t_i = 1$ for all ads, this problem reduces to the well-studied MULTIPLE KNAPSACK Problem with identical bin-capacities.

It should be observed that the AD PLACEMENT problem, as described at the beginning of this section, is one of *high-multiplicity* since the set of T banners, which are identical,

25

may be considered as one banner-type. This is also true of the jobs (ads), where identical copies of each job (ad) are encoded compactly. Hence T and the number of copies of each ad i, t_i, may be encoded in binary format, and therefore, a polynomial time algorithm for the problem would have to be polynomial in N, $\log T$, $\log c$ and $\log W$, where $c = \max_i\{c_i\}$. Recent papers in scheduling, including on the AD PLACEMENT problem, have addressed this question of high multiplicity encoding of the input and provide polynomial time algorithms even under such an encoding (see Clifford and Posner [18]). We will also present an efficient approximation scheme in Section 1.7.1 for a related variant of this problem, even when the problem input is specified using high-multiplicity encoding.

While the AD PLACEMENT problem described so far assumes a fixed height for all the ads, the AD PLACEMENT problem might also be studied in a more general setting, wherein ads have both height and width specifications and we require to place these rectangular ads in a rectangular banner display. We refer to this problem as the 2-DIMENSIONAL AD PLACEMENT problem. This problem was first studied by Adler et al. [4], however under very restrictive assumptions. In the special case of the problem when $T = 1$, this problem reduces to another problem of recent interest, namely the RECTANGLE PACKING problem [58, 97]. In this chapter, we also study a relaxation of the 2-DIMENSIONAL AD PLACEMENT problem and present some results for it.

1.5. Literature Survey

Adler et al. [4] were the first to study the algorithmic questions of advertisement scheduling and its related optimization variants, under both static (*offline*) and dynamic (*online*) settings. They assumed that the pricing scheme (profit) of the prospective advertisements was proportional to the space utilized by them. In other words, they assumed that $c_i = w_i$ in their model. Under this restriction, they studied the AD PLACEMENT problem in the form described in the previous section and showed it to be NP-hard via a reduction from the PARTITION problem. They proposed a 2-approximation algorithm for the problem, under the restriction that the ad sizes are divisible which however, performs arbitrarily badly if

ads sizes are not divisible. Apart from this problem, they also looked at the ad scheduling problem with ad geometries of higher dimensions under the same settings as the previous problem, and showed that any algorithm for the 1-dimensional case could be generalized to that of higher dimensions. For the online-version of the problem under the size-dependent pricing scheme and divisible size setting, they developed an online algorithm that is optimal. In many ways, their work was the main motivation for much of the subsequent work in ad scheduling.

Following the work by Adler et al. [4], the AD PLACEMENT problem under a size-dependent pricing scheme ($c_i = w_i$) was studied over a series of papers by Kumar, Sriskandarajah and their coauthors [57, 27, 26, 28], who referred to the problem as the $MAXSPACE$ problem. These papers developed a $\min\{10/3, 4L/(L + 1)\}$ for this $MAXSPACE$ problem and presented improved bounds for other special cases of the AD PLACEMENT problem. Amiri and Menon [7] performed a computational study of a Lagrangean-decomposition based algorithm to solve the $MAXSPACE$ problem. More recently, Kumar et al. [69] also present computational results for a hybrid genetic algorithm that they propose to solve the $MAXSPACE$ problem.

The AD PLACEMENT problem in its most general form, with the profit associated with each ad i, c_i, independent of its size, w_i, was studied by Freund and Naor [40]. They proposed a $(3 + \epsilon)$-approximation scheme for the AD PLACEMENT problem, which is polynomial even under the high-multiplicity encoding format. Their result even improved upon the previous best-known bound of Kumar et al. [28] for the $MAXSPACE$ problem, which is the special case of the AD PLACEMENT problem where $c_i = w_i$.

Another optimization variant of the ad scheduling problem, namely the problem of minimizing the width of a banner, given a set of ads to be necessarily allocated, was also studied by Adler et al. [4] and Kumar et al. [28], who referred to it as the $MINSPACE$ problem. In the case where $t_i = 1$ for all ads, this problem reduces to the classical NP-hard scheduling problem of minimizing the maximum completion time (makespan) over identical parallel machines, $P||C_{max}$, for which Graham [47] had developed a 4/3-approximation algorithm.

27

Kumar et al. [28] presented an LP-based 2-approximation algorithm and a 3/2-approximation algorithm for the $MINSPACE$ problem. This result was subsequently improved upon by Dean and Goemans [30] to 4/3, using a generalization of Graham's algorithm, whose running time is even strongly polynomial, and does not depend on the number of banners, T. Thus, the algorithm of Dean and Goemans [30] is polynomial even under high-multiplicity encoding. They also develop a polynomial time approximation scheme for the $MINSPACE$ problem, whose running time is polynomial in T.

As observed earlier, The AD PLACEMENT problem has a structure similar to that of a large number of scheduling problems, involving scheduling of parallel jobs on a set of identical processors. A very special case of the AD PLACEMENT problem is when $t_i = 1$ for all the ads, in which case the problem reduces to the MULTIPLE KNAPSACK problem with identical bin capacities. Another problem that is closely related to the AD PLACEMENT problem, and a relaxation that we study, is a weighted throughput maximization problem over multiple machines in real-time scheduling. In this problem, n jobs need to be scheduled on k identical parallel machines. Each job has associated with it, a release date, a deadline, an associated weight, and a processing time. The objective of the scheduler is to schedule jobs on machines, so as to maximize the weight of all feasibly scheduled jobs. Bar-Noy et al. [8] studied this problem, which may be represented as $P|r_i, d_i| \sum w_i(1 - U_i)$ in scheduling notation, and developed an LP-based approximation algorithm with ratios of $\frac{(1+1/k)^k}{(1+1/k)^k-1}$ for polynomially bounded integral input, and $\frac{(1+1/2k)^k}{(1+1/2k)^k-1}$ for arbitrary input, where k is the number of identical machines. It may be noted that by setting $r_i = 0$ and $d_i = d$ for all jobs, the throughput maximization problem reduces to the MULTIPLE KNAPSACK problem.

1.6. Problem Formulation and Complexity

The AD PLACEMENT problem described formally in Section 1.4 may also be formulated as an integer program. Let x_i denote a binary variable which indicates whether ad i belongs to a feasible solution, and the binary variable y_{ij} indicates whether the i^{th} ad is assigned to the j^{th} time banner. The AD PLACEMENT problem may now be rewritten as follows:

28

$$\text{maximize} \quad \sum_{i=1}^{N} c_i x_i$$

$$\text{subject to} \quad \sum_{i=1}^{N} w_i y_{ij} \leq W \quad j = 1, 2, \ldots, T, \tag{1.1a}$$

$$\sum_{j=1}^{T} y_{ij} = t_i x_i \quad i = 1, 2, \ldots, N, \tag{1.1b}$$

$$y_{ij} \in \{0, 1\} \quad i = 1, 2, \ldots, N, j = 1, 2, \ldots, T,$$

$$x_i \in \{0, 1\} \quad i = 1, 2, \ldots, N.$$

In this integer program, The set of constraints, (1.1a), ensure that the sum of the widths of the ads assigned to the j^{th} banner does not exceed W, the capacity of the j^{th} banner. The second set of constraints, (1.1b), ensure that in order for the i^{th} ad to belong to a feasible solution, all t_i copies of it must be assigned among the T banners. The 0-1 integrality constraints on y_{ij} ensure that no two copies of the same ad are assigned to the same banner. Finally, integrality constraints on x_i ensure that no copy of the i^{th} ad is assigned to any bin, unless there is feasible assignment for all t_i copies of ad i.

As discussed earlier, this problem is closely related to the MULTIPLE KNAPSACK problem and thus, we may interpret our problem as a generalization of the latter. Noting this similarity, in the rest of the paper we will use the words 'banner' and 'bin' interchangeably, and so also 'ad' and 'job'.

We now proceed to prove that the decision version of the AD PLACEMENT problem is strongly NP-complete, by a reduction from the 3-PARTITION problem. Note that this does not directly follow from the strong NP-hardness of the MULTIPLE KNAPSACK problem, since in that problem, the bin capacities need not be equal for all bins.

3-PARTITION

Instance: Nonnegative integers $a_1, a_2, \ldots, a_{3n}$; $\sum_{i=1}^{3n} a_i = nb$.

Question: Is there a partition of the integers into n sets $S_1, S_2, \ldots, S_n$ of 3 elements each, such that $\sum_{a_j \in S_i} a_j = b$, for $i = 1, 2, \ldots, n$?

Garey et al. [44] showed that 3-PARTITION is strongly NP-complete.

Theorem 1.1. AD PLACEMENT *is strongly NP-hard.*

Proof. Given an arbitrary instance of 3-PARTITION, we construct an associated instance of the decision version of AD PLACEMENT, which is polynomial in the size of the given instance. We will then show that the instance of 3-PARTITION is a "yes" instance if and only if the associated decision version of AD PLACEMENT is a "yes" instance.

Let us define $K = nb$, i.e. $K = \sum_{i=1}^{3n} a_i$. Consider the following associated decision instance of the AD PLACEMENT problem:

Instance: Ads $A_1, A_2, \ldots, A_{3n}$, each A_i specified by $(1, 1, K + a_i)$, for $i = 1, \ldots, 3n$; a banner of size $3K + b$; n time slots.

Question: Does there exist a feasible allocation of ads with profit at least $3n$?

Clearly the above instance of AD PLACEMENT is polynomial in the size of the corresponding instance of 3-PARTITION, since K is polynomial in the size of the 3-PARTITION instance.

Now suppose the instance of 3-PARTITION is a "yes" instance. In other words, there exists a partition of the $3n$ elements into $S_1, \ldots, S_n$ such that each set is of size 3 and the sum of the elements in each set is b. Consider the following allocation of ads. For each element $a_i \in S_j$, assign ad A_i to the j^{th} time slot banner. Observe that:

(a) Each ad A_i is assigned to some banner, since each corresponding element a_i is in some set S_j.

(b) For each banner j, exactly 3 ads are assigned to it, namely those ads A_i corresponding to the elements a_i in S_j. Also, since $\sum_{a_i \in S_j} a_i = b$, so $\sum_{A_i \in j} w(A_i) = 3K + \sum_{a_i \in S_j} a_i = 3K + b$.

(c) The allocation of ads is feasible, and so the profit obtained from this allocation is $3n$. Hence the corresponding AD PLACEMENT instance is also a "yes" instance.

Conversely, suppose the AD PLACEMENT instance is a "yes" instance. Then, since we have only $3n$ ads each with profit 1, all the $3n$ ads must be present in the feasible allocation. Also, note that since the banner size is $3K + b$, by definition of K, each time slot banner cannot accommodate more than 3 ads. Hence, exactly 3 ads must be allocated in a feasible manner to each time slot banner in the given "yes" instance. Now for each

ad A_i in banner j, assign element a_i to set S_j. Since, $\sum_{a_i \in S_j} a_i \leq b$ by this construction, and $\sum_{j=1}^{n} \sum_{a_i \in S_j} a_i = \sum_{i=1}^{3n} a_i = nb$, it follows that $\sum_{a_i \in S_j} a_i = b$. Hence, we have that the 3-PARTITION instance is also a "yes" instance. ∎

The above result implies that the AD PLACEMENT problem does not admit a fully polynomial time approximation scheme, unless P=NP. So, the best we can hope to achieve for this problem is a PTAS.

1.7. A Relaxation of the Ad Placement Problem

Freund and Naor [40] developed a $(3 + \epsilon)$-approximation algorithm for the AD PLACEMENT problem by studying a relaxation of the IP described in Section 1.6. Consider the following related KNAPSACK problem instance: Corresponding to each set of ads $A_i = (c_i, t_i, w_i)$, introduce a job i of width $t_i \cdot w_i$ and profit c_i. The problem is to find a packing of jobs of maximum profit in a knapsack of width $T \cdot W$. Clearly, any feasible solution to the AD PLACEMENT problem yields a feasible solution with the same profit for the above instance of the KNAPSACK problem, by concatenating the contents of all the T bins of the AD PLACEMENT instance solution. Hence, the optimal solution of the above instance of the KNAPSACK problem is an upper bound on the optimal profit of the AD PLACEMENT problem.

Thereafter, using the solution of the FPTAS for this knapsack problem instance, Freund and Naor [40] determined the set of corresponding prospective ads, S, that they would attempt to pack feasibly among the T banners. Clearly, if all the ads in S can be packed feasibly among the T banners, then this would be an $(1 + \epsilon)$-approximate solution to the AD PLACEMENT problem. Using an $SSLF$ (Smallest Size first Least Full bin first) assignment policy for the ads amongst the banners, they partitioned S into three sets of ads S_1, S_2 and S_3, for each of which they showed feasible packings among the T banners. By selecting the subset of ads among S_1, S_2, S_3 with the highest profit, they showed that their algorithm was in fact a $(3 + \epsilon)$-approximation scheme and that their analysis was tight. The reader is referred to [40] for details of their $(3 + \epsilon)$-approximation result.

In order to develop a better approximation algorithm for the AD PLACEMENT problem, we will study an alternate relaxation of the above problem, where we relax integrality constraints on x_i. In other words, $0 \leq x_i \leq 1$, for $i = 1, 2, \ldots, N$. This ILP relaxation then appears as follows:

$$\text{maximize} \quad \sum_{i=1}^{N} c_i \cdot x_i$$

$$\text{subject to} \quad \sum_{i=1}^{N} w_i \cdot y_{ij} \leq W \quad j = 1, 2, \ldots, T, \tag{1.2a}$$

$$\sum_{j=1}^{T} y_{ij} = t_i \cdot x_i \quad i = 1, 2, \ldots, N, \tag{1.2b}$$

$$y_{ij} \in \{0, 1\} \quad i = 1, 2, \ldots, N, j = 1, 2, \ldots, T,$$

$$0 \leq x_i \leq 1 \quad i = 1, 2, \ldots, N.$$

In the above relaxation of the AD PLACEMENT problem, it is possible to eliminate all the x_i variables, which in turn yields the following reduced integer program, (1.3):

$$\text{maximize} \quad \sum_{j=1}^{T} \sum_{i=1}^{N} \frac{c_i}{t_i} \cdot y_{ij} \tag{1.3}$$

$$\text{subject to} \quad \sum_{i=1}^{N} w_i \cdot y_{ij} \leq W \quad j = 1, 2, \ldots, T,$$

$$\sum_{j=1}^{T} y_{ij} \leq t_i \quad i = 1, 2, \ldots, N,$$

$$y_{ij} \in \{0, 1\} \quad i = 1, 2, \ldots, N, j = 1, 2, \ldots, T.$$

The above formulation has a very intuitive interpretation, and may be thought of as an alternate banner advertisement placement problem: we again have N sets of possible ads, with each ad i having t_i copies of weight w_i and profit $\frac{c_i}{t_i}$ each. We want a feasible assignment of these copies of ads to T banners of size W each, subject to the constraint that no two copies of the same ad can be assigned the same banner. Our objective is again to maximize the profit obtained from a feasible assignment. The key aspect that makes this problem (1.3) different from the AD PLACEMENT problem is that in this problem we can assign fewer than t_i copies of any ad i in a feasible assignment.

While (1.3) is a relaxation of the AD PLACEMENT problem, it turns out that the problem still remains strongly NP-hard.

Lemma 1.2. *The relaxed* AD PLACEMENT *problem (1.3) remains strongly NP-hard.*

Proof. The proof again follows from a reduction of the 3-PARTITION problem to 1.3, by the same construction used in Theorem 1.1. ∎

Although the relaxed problem, (1.3), is harder than the other KNAPSACK relaxation of the AD PLACEMENT problem considered by Freund and Naor [40], it still possesses an interesting structure. We therefore attempt to devise good approximation algorithms for this problem. We begin by first observing that any LP-relaxation based algorithm for (1.3) cannot provide an approximation guarantee better than 2. In other words, possibly an alternate formulation of the problem and its relaxation might in fact provide the basis for a better-than-2 LP-relaxation based algorithm, but this does not hold for the simple LP-relaxation of (1.3).

Lemma 1.3. *The integrality gap of (1.3) with respect to the standard LP relaxation is at least 2.*

Proof. We prove the above claim by giving an example of the problem where the LP/IP gap is 2. Consider the following instance of (1.3): We are given two ads, each with an exposure requirement of T. Both ads have a size of $W/2 + \epsilon$, and profit c each. Also, banners have size W and we have T banners in which to assign the ads.

Clearly, an optimal solution of the LP relaxation of the problem would be:

$$y_{ij} = W/(W + 2\epsilon), \quad i = 1, 2 \quad j = 1, \ldots, T,$$

yielding an optimal profit of $2cW/(W + 2\epsilon)$. It is clear however, that in an optimal solution to (1.3), only one copy of either of the two ads can be feasibly assigned to a bin, yielding a profit of profit of c. Hence, for this example:

$$OPT(LP)/OPT(IP) = 2W/(W + 2\epsilon) \approx 2. \qquad \blacksquare$$

We now present a greedy approximation algorithm for the relaxed AD PLACEMENT problem and show that it in fact guarantees a better-than-2 approximation result for the problem.

1.7.1 A Greedy Algorithm

GREEDY ALGORITHM

Step 1: Set $i = 1$; let L represent the set of indices of ads, not all of whose copies have been to assigned to bins. Initially, $L = \{1, 2, \ldots, N\}$.

Step 2: For bin i, solve the KNAPSACK problem, with the set of prospective jobs consisting of exactly one representative copy from each ad whose index is in L, and job i having a profit $\frac{c_i}{t_i}$. Let S_i represent the set of ads whose copies have been selected to be placed in the i^{th} bin.

Step 3: If a copy of ad j is in S_i, then $t_j := t_j - 1$;

Step 4: If $(t_j = 0)$ for job j, then $L := L \setminus \{j\}$.

Step 5: Set $i := i + 1$. If $i \leq T$, then goto **Step 2**.

Informally speaking, a greedy approach is employed by the above algorithm, where the KNAPSACK problem is solved T times in order to determine the contents of the T bins. Since while solving any single KNAPSACK problem, we consider only one representative copy of each ad, we always ensure that no more than one copy of the same ad in present in any bin. Thus we ensure that the solution generated by the above algorithm is indeed a feasible solution to the problem (1.3).

The running time of this greedy algorithm, however, is pseudopolynomial in the size of the input, as it solves the KNAPSACK problem $O(T)$ times, whereas the input may be represented compactly in $O(\log T)$ using high-multiplicity encoding. We now proceed to improve the computational complexity of the greedy algorithm, in order that its running time is polynomial in $\log T$. We describe a modified greedy algorithm below and show how it has an improved running time without affecting the performance guarantee of the original greedy algorithm. In the new algorithm, if S_i is not explicitly computed, it means that

34

$S_i = S_{i-1}$.

MODIFIED GREEDY ALGORITHM

Step 1: Set $i = 1$; let L represent the set of indices of ads, not all of whose copies have been to assigned to bins. Initially, $L = \{1, 2, \ldots, N\}$.

Step 2: For bin i, solve the KNAPSACK problem, with the set of prospective jobs consisting of exactly one representative copy from each ad whose index is in L, and job i having a profit $\frac{c_i}{t_i}$. Let S_i represent the set of ads whose copies have been selected to be placed in the i^{th} bin.

Step 3: Find $\bar{t} = \min_{j \in S_i}\{t_j\}$ and set $t_j = t_j - \bar{t} \quad \forall j \in S_i$;

Step 4: If $(t_j = 0)$ for a job j, then $L := L \setminus \{j\}$.

Step 5: Set $i := i + \bar{t}$. If $i \leq T$, then goto **Step 2**.

To prove that the equivalence of the solution of the greedy algorithm and the modified greedy algorithm, we make the following important observation:

Lemma 1.4. *In the original greedy algorithm, an α-approximate solution S_i calculated for the i^{th} bin remains α-approximate with respect to profits for the following $\bar{t} = \min_{j \in S_i}\{t_j\}$ iterations.*

Proof. Follows from the fact that all the elements in S_i remain in L for the next $\bar{t}$ iterations, after which the element in S_i with the least number of copies available may get eliminated from L. Moreover, in each successive iteration, the value of the corresponding optimal solution for the i^{th} bin is nonincreasing, while the objective value of the approximate solution remains the same. ∎

Having made this observation, the equivalence of the two greedy algorithms, in terms of the profit they compute, is obvious. All that remains to be shown is that the modified greedy algorithm has a superior running time.

Lemma 1.5. *The number of iterations of the modified greedy algorithm is bounded by N, the number of prospective ads. As a consequence, the running time of the modified greedy*

algorithm, using a pseudopolynomial time algorithm to compute the optimal KNAPSACK *solution, is* $O(N^2W)$.

Proof. In the modified greedy algorithm, an optimal bin configuration S_j is computed only if at least one prospective ad i is removed from L. Hence the number of iterations of the greedy algorithm, which is the number of times an optimal bin configuration is computed, is bounded by N. Each optimal computation of a bin configuration takes at most $O(NW)$ time, using the pseudopolynomial algorithm for the KNAPSACK problem. ∎

Whereas in the original greedy algorithm the number of iterations was always T, we are able to reduce the number of iterations now to at most N. Hence the dependence of this new greedy algorithm is only polynomial in the encoding of T. Henceforth, while for the sake of explanation we will use the original greedy algorithm, the actual greedy algorithm that will be implemented will be the modified greedy algorithm.

1.7.2 Analysis of the Greedy Algorithm for Ad Placement

We begin by showing that the relaxed AD PLACEMENT problem (1.3) problem is in fact an instance of the more general problem of maximizing a *normalized, nondecreasing, submodular function* over a *uniform matroid*. The reader may recall the definitions of these terms from Section 1.2.

Lemma 1.6. *The relaxed* AD PLACEMENT *problem, (1.3), is an instance of maximizing a normalized, nondecreasing, submodular function over a uniform matroid.*

Proof. Let S represent a set of ads, such that the sum of the size of one copy of each ad does not exceed W (i.e. $\sum_{A_j \in S} w_j \leq W$). Let s represent the "type" of a bin containing exactly one copy of each ad in S. For ease of representation, we say that ad $i \in s$ whenever $A_i \in S$. Notice that any feasible assignment of the relaxed AD PLACEMENT problem (1.3) can contain at most $\min_{i \in s}\{t_i\}$ bins of the bin-type s corresponding to S. Now, construct a ground set E of bins, containing exactly $\min_{i \in s}\{t_i\}$ bins of type s, for each possible bin-type, s. Note that the size of E may potentially be exponential in the size of the input.

Since in the relaxed AD PLACEMENT problem, we can use at most T bins, it follows that a feasible solution to (1.3) contains at most T elements from E, with an additional restriction that the ad assignment is feasible, in that no ad i is assigned more than t_i times in the T bins. For the time being, we ignore this restriction, as we shall capture it in the definition of f by not including the profits for the copies of an ad i beyond the first t_i copies. Now it follows that the set of feasible T-bin configurations of (1.3), say $\mathcal{F}$, is a uniform matroid, since $\mathcal{F} := \{F \subseteq E : |F| \leq T\}$.

Define a function $f : E \to \mathbb{R}_+$, with $f(\emptyset) = 0$ as follows:

$$f(e) = \sum_{i \in e} \frac{c_i}{t_i} \qquad \text{for } e \in E.$$

In other words, f is a measure of the profit of a feasible bin-type. It is obtained by summing the profits of individual ad copies i in the bin-type e. For any subset $A \subseteq E$, let n_i^A represent the number of copies of ad i in the $|A|$-bin configuration given by A. Clearly, $n_i^A \leq |A|$. We define $f(A)$ as:

$$f(A) = \sum_{i=1}^{N} \frac{c_i}{t_i} \cdot \min\{n_i^A, t_i\}$$

Thus, we ensure that for any ad i, only the first t_i copies are allocated a profit of $\frac{c_i}{t_i}$ each, in any bin configuration A. Observe now that:

- f is *nondecreasing*: Consider any set $A \subseteq E$ and add to it another element, say $e \in E \setminus A$. Now,

$$f(A \cup \{e\}) = \sum_{i=1}^{N} \frac{c_i}{t_i} \cdot \min\{n_i^{A \cup \{e\}}, t_i\} \geq \sum_{i=1}^{N} \frac{c_i}{t_i} \cdot \min\{n_i^A, t_i\} = f(A).$$

The inequality follows from the fact that $\min\{n_i^{A \cup \{e\}}, t_i\} \geq \min\{n_i^A, t_i\}$.

- f is *submodular*: Suppose $A \subseteq B \subseteq E$ and $e \in E \setminus B$. We define a new binary variable $v_i(S)$ corresponding to ad i and a set $S \subseteq E$ to be: $v_i(S) = 1$ if $n_i^S \geq t_i$ and $v_i(S) = 0$ otherwise. Notice that $v_i(B) \geq v_i(A)$ for any ad i, since $A \subseteq B$. Now, clearly:

$$f(A\cup\{e\})-f(A) = \sum_{i\in e}\tfrac{c_i}{t_i}\cdot(n_i^{\{e\}}-v_i(A)) \geq \sum_{i\in e}\tfrac{c_i}{t_i}\cdot(n_i^{\{e\}}-v_i(B)) = f(B\cup\{e\})-f(B).$$

So we have proved that f is normalized, nondecreasing, submodular. It may be noted that the function f that we have defined measures exactly the objective function value of the relaxed AD PLACEMENT problem. ∎

Suppose that we use an α-approximation algorithm for KNAPSACK in the modified greedy algorithm presented earlier. Clearly, given the existence of an FPTAS for it, we know that we could use a good enough approximation factor, α. The greedy algorithm proposed for (1.3) in Section 1.7.1 determines a bin e to add at each iteration as the bin e that approximately maximizes the incremental function value, $\rho_e(S) = f(S\cup\{e\})-f(S)$, at each iteration using an α-approximate algorithm for KNAPSACK over the remaining set of ads, L. Note that the set S in the above statement is the bin configuration selected by the greedy algorithm up to the iteration in question. As it turns out, the greedy algorithm presented for the relaxed AD PLACEMENT problem may be generalized to a greedy algorithm used for maximizing nondecreasing submodular functions over a uniform matroid, and for whose performance, we prove a more general approximation result in Chapter 2 (refer Theorem 2.3).

Our main result for the performance of the modified greedy algorithm in solving (1.3) is as follows:

Theorem 1.7. *The modified greedy algorithm, using an α-approximation algorithm to solve the KNAPSACK problem, is in fact an $\frac{e^{\frac{1}{\alpha}}}{e^{\frac{1}{\alpha}}-1}$-approximation algorithm for the relaxed AD PLACEMENT problem, (1.3).*

Proof. Follows as a consequence of Theorem 2.3, where we show that a more general form of the modified greedy algorithm is in fact an $\frac{e^{\frac{1}{\alpha}}}{e^{\frac{1}{\alpha}}-1}$-approximation algorithm for the problem of maximizing a nondecreasing submodular function over a uniform matroid, of which the relaxed AD PLACEMENT problem is a special case. ∎

Now, using that fact that there is an $\alpha = (1+\epsilon)$-approximation scheme for the KNAPSACK problem, it follows that:

Corollary 1.8. *The modified greedy algorithm, with the FPTAS for the* KNAPSACK *problem as a subroutine, is an* $(\frac{e}{e-1} + \epsilon)$*- approximation scheme, whose running time is* $O(\frac{N^4}{\epsilon})$.

Chekuri and Khanna [15] showed the same result for the performance of the greedy algorithm in solving the MULTIPLE KNAPSACK problem with identical bin capacities, a special case of the relaxed AD PLACEMENT problem.

Another special case of the AD PLACEMENT problem is where $t_i = t$ for $i = 1, \ldots, N$ and the number of banners, T, is an integral multiple of t. This special case of the problem was considered by Dawande et al. [28], under the assumption that $c_i = w_i$ and they proposed a 2-approximation algorithm for it, that was tight. It follows from the description of the modified greedy algorithm that it outputs a feasible solution of this special version of the AD PLACEMENT problem. Hence, we may make the following observation that improves upon their 2-approximation result:

Observation 1.9. *For the special case of the* AD PLACEMENT *problem where* $t_i = t$ *for* $i = 1, \ldots, N$ *and the number of banners, T, is an integral multiple of t, the modified greedy algorithm with the FPTAS for the* KNAPSACK *problem as a subroutine, is an* $(\frac{e}{e-1} + \epsilon)$*-approximation scheme, even if* $c_i \neq w_i$ *for* $i = 1, \ldots, N$.

1.8. The 2-Dimensional Ad Placement Problem

The 2-DIMENSIONAL AD PLACEMENT problem is an extension of the AD PLACEMENT problem, wherein we now consider rectangular ads that need to be placed in a rectangular display area, without constraining ads to be of a specified height. In this problem, ads may be specified by a 4-tuple, (l_i, w_i, c_i, t_i), where l_i and w_i represent the height and width of the ad, t_i represents the exposure of the ad, and c_i is the profit associated with a feasible assignment of the ad. Also, we are given T rectangular banners of height L and width W each, in which we have to determine a *feasible* allocation of the ads, so that:

(a) no two copies of the same ad are assigned to the same banner,

(b) any ad, A_i, in the allocation, appears in at least t_i of the T banners, and

(b) no two ads ever overlap with each other, and ads may not be rotated.

As the ad scheduler, our objective is to maximize the profit of a feasible allocation of ads. The problem may be defined more formally as follows:

2-DIMENSIONAL AD PLACEMENT

Instance: Rectangular ads $A_1, A_2, \ldots, A_N$, each ad A_i specified by height l_i, width w_i, exposure requirement t_i and profit c_i, for $i = 1, \ldots, N$; a rectangular banner of height L and width W ; T time slots. $(T \geq t_i$, for $i = 1, \ldots, N)$

Task: Find a subset, $S \subseteq \{1, \ldots, n\}$ of ads such that all ads A_i for $i \in S$ may be assigned amongst the T banners *feasibly*, so as to maximize profit, $\sum_{i \in S} c_i$.

Since the AD PLACEMENT problem is a special case of the above problem, the strong NP-hardness of this problem follows. The AD PLACEMENT problem in higher dimensions was first studied by Adler et al. [4], who proposed an algorithm that constructs a feasible solution to the higher dimensional problem, given an algorithm that constructs a feasible solution for the AD PLACEMENT problem. However, this algorithm only worked under a setting where ad dimensions were *divisible*. In this section, we will state some results for a relaxation of the 2-DIMENSIONAL AD PLACEMENT problem, analogous to the relaxation of the AD PLACEMENT problem (1.3) that we studied in Section 1.7.

Consider the following relaxation of the 2-DIMENSIONAL AD PLACEMENT problem: We again have N sets of possible ads, with each ad A_i having t_i copies of size (l_i, w_i) and profit $\frac{c_i}{t_i}$ each. We want a feasible assignment of these copies of ads to T banners of size (L, W) each, subject to the constraint that no two copies of the same ad can be assigned the same banner. Our objective is again to the maximize the profit obtained from a feasible assignment. This problem is different from the 2-DIMENSIONAL AD PLACEMENT problem because in this problem we may assign fewer than t_i copies of any ad i in a feasible assignment and make the corresponding fraction of a profit.

Consider the case when $t_i = T = 1$ for all $i = 1, \ldots, N$. In this case, both the 2-DIMENSIONAL AD PLACEMENT problem and its relaxation reduce to the RECTANGLE PACKING problem that was mentioned in Section 1.2.3. It may easily be seen that if we

40

model a set E consisting of feasible 1-banner display configurations, then through a transformation similar to that used for the relaxed AD PLACEMENT problem in Lemma 1.6, the relaxation of the 2-DIMENSIONAL AD PLACEMENT problem also reduces to the problem of maximizing a submodular function over a uniform matroid. Moreover, one might think of an extension of the greedy algorithm presented in Section 1.7.1, wherein, instead of solving the KNAPSACK problem to determine a configuration for each banner, the greedy algorithm solves a RECTANGLE PACKING problem to determine the configuration of each banner. This greedy algorithm for the relaxed 2-DIMENSIONAL AD PLACEMENT problem may be stated as follows:

GREEDY ALGORITHM FOR RELAXED 2-DIMENSIONAL AD PLACEMENT

Step 1: Set $i = 1$; let L represent the set of indices of ads, not all of whose copies have been to assigned to bins. Initially, $L = \{1, 2, \ldots, N\}$.

Step 2: For bin i, solve the RECTANGLE PACKING problem, with the set of prospective jobs consisting of exactly one representative copy from each ad whose index is in L, and job i having a profit $\frac{c_i}{t_i}$. Let S_i represent the set of ads whose copies have been selected to be placed in the i^{th} bin.

Step 3: If a copy of ad j is in S_i, then $t_j := t_j - 1$;

Step 4: If $(t_j = 0)$ for job j, then $L := L \setminus \{j\}$.

Step 5: Set $i := i + 1$. If $i \leq T$, then goto **Step 2**.

As with the greedy algorithm for the relaxed AD PLACEMENT problem, one might also modify the greedy algorithm presented above to ensure that it is still polynomial under high-multiplicity encoding. Using the 2-approximation algorithm of Jansen and Zhang [58] for the RECTANGLE PACKING problem, our result for the performance of the greedy algorithm in solving the relaxed 2-DIMENSIONAL AD PLACEMENT problem is that:

Theorem 1.10. *The greedy algorithm, using the $(2 + \epsilon)$-approximation algorithm for the* RECTANGLE PACKING *problem, is an $(\frac{\sqrt{e}}{\sqrt{e}-1} + \epsilon)$-approximation scheme for the relaxed 2-* DIMENSIONAL AD PLACEMENT *problem.*

Proof. Follows as a consequence of Theorem 2.3, where we show that a more general form of the greedy algorithm presented above is in fact an $\frac{e^{\frac{1}{\alpha}}}{e^{\frac{1}{\alpha}}-1}$-approximation algorithm for the problem of maximizing a nondecreasing submodular function over a uniform matroid, of which the relaxed AD PLACEMENT problem is a special case. Moreover the role of the α-incremental oracle discussed in the more general greedy algorithm is fulfilled by the $(2+\epsilon)$-approximation scheme for the RECTANGLE PACKING problem, yielding that $\alpha = 2 + \epsilon$. ∎

1.9. Search-based Advertising: The AdWords Assignment Problem

Search-based advertising is, increasingly, the most popular form of advertising on the Internet, and a significant source of revenue for search portals such as Google, Yahoo, and MSN, to name a few. In this paradigm, portals solicit advertisements for particular *keywords*, called "AdWords" in the case of Google. The phrase, "linear programming," would be an example of one such keyword, with which advertisers might like to associate their advertisement. When a user of a search portal types in a *query*, it is matched with a corresponding keyword and the associated advertisements are displayed on the search results page. Hence, advertisers would have different *valuations* for their advertisement being associated with different keywords, which in turn, would depend on which queries match each keyword.

If a search firm is aware of advertisers' private valuations and, therefore, their willingness to pay for each keyword, then one might think of framing the search firm's problem of assigning advertisements to keywords to maximize revenue as the following optimization problem proposed by Fleischer et al. [39]:

ADWORDS ASSIGNMENT [39]

Instance: A set of n bidders with a budget B_i for each bidder i; a rectangular ad A_i of length l_i and width w_i that bidder i would like to advertise; a set of m AdWords (keywords), each AdWord j with a rectangular display area of length L_j and width W_j; bidder i has a maximum willingness to pay, v_{ij}, for having its ad associated with AdWord j.

Task: Find a feasible assignment of AdWords, S_i, to bidder i so as to maximize total revenue, where the revenue obtained from bidder i is $\min(B_i, \sum_{j \in S_i} v_{ij})$. In a feasible assignment, all ads assigned to AdWord j must be feasibly displayed in the rectangular display without having to rotate any of the ads.

It turns out this ADWORDS ASSIGNMENT problem also has an underlying structure very similar to the relaxed AD PLACEMENT problem described earlier. It is in fact a special case of maximizing a nondecreasing submodular function over a *partition matroid* (recall from Section 1.2). We present this transformation below.

Lemma 1.11. *The* ADWORDS ASSIGNMENT *problem is a special case of maximizing a nondecreasing submodular function over a partition matroid.*

Proof. Consider an underlying ground set $E = \cup_{j=1}^{m} E_j$, where each element $e_S \in E_j$ corresponds to a feasible assignment of ads, S, that may be accommodated in the rectangular display corresponding to AdWord j. A feasible solution to the ADWORDS ASSIGNMENT problem would constrain that at most one element may be picked of each type, E_j, thereby defining a partition matroid on E.

For any subset, $F \subseteq E$, suppose that $F_j = F \cap E_j$. Furthermore, let $F_j^i = \{e_S \in F_j : \text{ad } i \in S\}$. Define a function, f, on a subset $F \subseteq E$ as follows:

$$f(F) = \sum_{i=1}^{n} \min(B_i, \sum_{j=1}^{m} v_{ij}|F_j^i|)$$

It is not hard to verify that f is exactly the objective function of ADWORDS ASSIGNMENT over all feasible sets in the partition matroid, and therefore over all feasible solutions of AD-

43

WORDS ASSIGNMENT. Moreover, observe that f is indeed nondecreasing and has decreasing marginal values, by construction. Hence the claim holds. ∎

For the problem of maximizing a nondecreasing, submodular function over a partition matroid, we later show in Chapter 2, that even a simple variant of the greedy algorithm, namely the *locally greedy* algorithm proposed by Fisher et al. [38], is an $(\alpha+1)$-approximation algorithm, given an α-approximation algorithm to solve a certain increment maximization problem. Based on this result, we are able to develop a polynomial-time $(3+\epsilon)$-approximation scheme for ADWORDS ASSIGNMENT (see Theorem 2.9). Consequently, we are able to improve on a previous $(3.16 + \epsilon)$-approximation result due to Fleischer et al. [39] for the problem. We defer the presentation of this improved approximation algorithm and the proof of its performance to Chapter 2.

In practice, however, it is hard for search firms to know the willingness to pay of advertisers and, consequently, search firms such as Google and Yahoo use a *generalized second price auction* mechanism to determine not only which ads are matched to a keyword, but also to determine the position of an ad on the search results page. Moreover, there is also an additional challenge specific to the context of Internet search: a user *query* to a search engine might correspond to more than one keyword, thereby creating the challenge of different keyword auctions competing against each other to appear of the search results page. This therefore raises a problem for the advertiser, to determine how much it must bid on each keyword so as to maximize its objective, whether the objective be clicks or "actions" (such as subscribing for a newsletter, or signing up for an account, etc.). Of course, the advertiser only has a limited budget to invest on such search-based advertisements. Feldman et al. [37] study this BUDGET OPTIMIZATION problem for the advertiser, given that a search firm uses a generalized second price auction mechanism. We refer the reader to their paper for further details about their model and results, but we do note that even if the objective of the advertiser is performance-based, with different weights associated for each *query*, the BUDGET OPTIMIZATION problem may in fact be framed as maximizing a nondecreasing submodular function over a submodular knapsack constraint. Thus the inherent underlying structure of

even this problem is closely related to that of the other optimization problems discussed in this chapter. We leave as an open question whether this more general view of the BUDGET OPTIMIZATION problem might yield improved approximation results.

1.10. Conclusion

In this chapter, we studied optimization problems that arise in the context of Internet advertising. Within the banner-based advertising paradigm, we considered the AD PLACEMENT problem and other related variants and developed approximation algorithms for some of them. More importantly, we observed a close relationship between assignment problems similar in spirit to the AD PLACEMENT problem, and submodular function maximization. Even in problems that arise in search-based advertising, such as the ADWORDS ASSIGNMENT problem, we observe a similar relationship. Building on these insights and observations, we address the more general problem of maximizing nondecreasing submodular set functions over special classes of constraints in the next chapter.

Chapter 2

Revisiting the Greedy Approach to Submodular Function Maximization

2.1. Introduction

Submodular set functions are widely used in the economics, operations research, and computer science literature to represent consumer valuations, since they capture the notion of *decreasing marginal utilities* (or alternatively, *economies of scale* in a cost framework). While these properties make submodular functions a suitable candidate of choice for objective functions, submodular objective functions also arise as a natural structural form in many classic discrete optimization settings, such as the MAX SAT problem in Boolean logic, the MAX CUT problem in graphs, and the MAXIMUM COVERAGE problem in location analysis, to name a few.

The role of submodularity in discrete optimization is akin to that of convex functions in continuous optimization, given their analogous prevalence, structural properties, and the tractability of solving minimization problems on both functions (Lovász [72], Fujishige [42]). Interestingly, submodular functions are also closely related to concave functions, and this raises the question of the tractability of maximizing submodular functions. However, since many NP-hard problems may be reduced to the problem of maximizing submodular func-

47

tions, it is unlikely that there exists a polynomial-time algorithm to solve this problem (unless P=NP). Consequently, a vast body of literature has focussed on developing efficient heuristics for various instances of this problem.

The greedy algorithm, that iteratively augments a current solution with an element of maximum incremental value, has been shown to be an effective heuristic in maximizing nondecreasing submodular functions over different constraint structures [19, 34, 38, 78, 80, 99, 107]. In most prior works, it was implicitly assumed that the greedy algorithm has access to an *incremental oracle* that, given a current solution, returns in polynomial time an element of highest incremental value to the current solution. However, it turns out that in some problems, determining an element with the best incremental profit may itself be an NP-hard problem, thus necessitating the use of only an *approximate* incremental oracle. In this chapter, we generalize the performance bounds of the greedy algorithm and an interesting related variant, the *locally greedy* algorithm (Fisher et al. [38]) for maximizing nondecreasing submodular set functions over various constraint structures, when the algorithm only has access to an approximate incremental oracle. Subsequently, we discuss how various results in the modern literature for problems that arise in the context of assignment problems, Internet advertising, wireless sensor networks, combinatorial auctions, and utility games among others, may be reinterpreted using these generalized performance bounds.

2.1.1 Preliminaries

A real-valued set function $f : 2^E \to \mathbb{R}$ is **normalized, nondecreasing** and **submodular** if it satisfies the following conditions, respectively:

(F0) $f(\emptyset) = 0$;

(F1) $f(A) \leq f(B)$ whenever $A \subseteq B \subseteq E$;

(F2) $f(A) + f(B) \geq f(A \cup B) + f(A \cap B)$ for all $A, B \subseteq E$, or equivalently:

(F2a) $f(A \cup \{e\}) - f(A) \geq f(B \cup \{e\}) - f(B)$ for all $A \subseteq B \subseteq E$ and $e \in E \setminus B$, or equivalently:

(F2b) $f(A \cup C) - f(A) \geq f(B \cup C) - f(B)$ for all $A \subseteq B \subseteq E$ and $C \subseteq E \setminus B$.

Henceforth, whenever we refer to submodular functions, we shall, in particular, imply normalized, nondecreasing, submodular functions. We also adopt the following notation: For any two sets $A, B \subseteq E$, we define the **marginal value** (incremental value) of set A to set B as

$$\rho_A(B) = f(A \cup B) - f(B).$$

Additionally, we will use the subscript e instead of $\{e\}$ whenever the context is clear. In particular, **(F2a)** can equivalently be written as $\rho_e(A) \geq \rho_e(B)$ for $A \subseteq B$.

A set system $(E, \mathcal{F})$, where E is a finite set and $\mathcal{F}$ is a collection of subsets of E, is an **independence system** if it satisfies the following properties:

(M1) $\emptyset \in \mathcal{F}$;

(M2) If $X \subseteq Y \in \mathcal{F}$ then $X \in \mathcal{F}$.

Furthermore, any set $X \in \mathcal{F}$ is called an **independent** set, whereas a set $Y \in 2^E \setminus \mathcal{F}$ is called a **dependent** set. A maximal independent set in $\mathcal{F}$ is called a **basis**.

An independence system $(E, \mathcal{F})$ is a **matroid** if it satisfies the additional property:

(M3) If $X, Y \in \mathcal{F}$ and $|X| > |Y|$, then there is an $x \in X \setminus Y$ with $Y \cup \{x\} \in \mathcal{F}$.

Matroids have the property that all bases in $\mathcal{F}$ have the same cardinality.

In this chapter, we will also focus our attention on the following special classes of matroids:

- **Uniform matroids**: E is a finite set, k is a positive integer, and

$$\mathcal{F} := \{F \subseteq E : |F| \leq k\}.$$

- **Partition matroids**: $E = \cup_{i=1}^{k} E_i$ is the disjoint union of k sets, $l_1, \ldots, l_k$ are positive integers, and

$$\mathcal{F} = \{F : F = \cup_{i=1}^{k} F_i \text{ where } F_i \subseteq E_i, |F_i| \leq l_i \text{ for } i = 1, \ldots, k\}.$$

- **Laminar matroids** (Calinescu et al. [14], Gabow and Kohno [43]): Let E be a finite set. A family of subsets $\mathcal{S} \subseteq 2^E$ is said to be a laminar family if for any two sets

$X, Y \in \mathcal{S}$, at least one of the three sets, $X \setminus Y$, $Y \setminus X$, $X \cap Y$ is empty. Let $\mathcal{S}$ be a laminar family of sets, and each set $S \in \mathcal{S}$ is associated with an integer value k_S. Then,

$$\mathcal{F} = \{F \subseteq E : |F \cap S| \leq k_S \text{ for each } S \in \mathcal{S}\}.$$

Starting with the seminal work of Edmonds [32], submodularity and matroids have received a lot of attention in the optimization community. The reader is referred to standard textbooks in combinatorial optimization [21, 67, 79, 93] for a detailed exposition on submodularity, matroids, and independence systems.

Since our focus is on developing approximation algorithms for a variety of problems, we must formalize the notion of an approximation algorithm. An α-**approximation algorithm** for a maximization problem $\mathcal{P}$ is a polynomial-time algorithm A for $\mathcal{P}$ such that

$$OPT(I) \leq \alpha \cdot A(I)$$

for all instances I of $\mathcal{P}$, where $OPT(I)$ and $A(I)$ are the optimal value and the objective value returned by the algorithm A for an instance I of $\mathcal{P}$. Observe that by this definition, it must be that $\alpha \geq 1$. A **fully polynomial-time approximation scheme** (FPTAS) provides, for every $\epsilon > 0$, a $(1+\epsilon)$-approximation algorithm whose running time is polynomial in both the size of the input and $1/\epsilon$. More generally, a **polynomial-time approximation scheme** (PTAS) provides a $(1+\epsilon)$ approximation algorithm whose running time is polynomial in the size of the input, for any constant ϵ.

Finally, an inequality that will be useful in some of the analysis presented in this chapter is the **power means inequality**. Given n nonnegative numbers $a_1, a_2, \ldots, a_n$, and two nonzero real numbers p, q with $p > q$, the inequality states that:

$$\left(\frac{a_1^p + a_2^p + \ldots + a_n^p}{n}\right)^{1/p} \geq \left(\frac{a_1^q + a_2^q + \ldots + a_n^q}{n}\right)^{1/q} \tag{2.1}$$

with equality holding if and only if $a_1 = a_2 = \ldots = a_n$.

2.1.2 Problem Description and Literature Survey

The problem we address may be stated as follows:

$$z_{opt} = \max\{f(S) : S \subseteq E, S \in \mathcal{F}\} \qquad (P)$$

where f is a normalized, nondecreasing, submodular function, and $(E, \mathcal{F})$ is, in general, an independence system. As stated earlier, our focus is on the performance of the greedy algorithm (and its variants), described below, in solving some special cases of this general problem.

STANDARD GREEDY ALGORITHM

Initialization: $S := \emptyset$, $E' := E$.

Incremental Oracle: Select an element $e^* \in E' \setminus S$ such that

$$e^* = \operatorname*{argmax}_{e \in E' \setminus S} \rho_e(S).$$

Admissibility Oracle: If $S \cup \{e^*\} \in \mathcal{F}$

$$\text{Then } S := S \cup \{e^*\}.$$

$$\text{Else } E' := E' \setminus \{e^*\}.$$

Loop back: While $E' \setminus S \neq \emptyset$ goto **Incremental Oracle**.

End

Informally stated, the greedy algorithm starts with an empty set, and in each iteration adds an element with highest marginal value to the solution using an *incremental oracle*, while ensuring independence of the resulting solution set using an *admissibility oracle* (also known as independence oracle). The algorithm continues as long as there remains an element which it has not previously considered.

A special case of the problem (P) is the maximization of a linear function over a matroid. For this problem, the greedy algorithm is known to be optimal (Rado [87], Edmonds [32]). Korte and Hausmann [66] studied the problem of maximizing a linear function over an independence system and present tight bounds (that are functions of the rank quotient) on

51

the performance of the greedy algorithm for this problem. Nemhauser et al. [80] considered the problem (P) over a uniform matroid and showed that greedy is a tight $(e/(e-1))$-approximation algorithm for this problem. In a companion paper, Fisher et al. [38] studied the problem (P) over a general independence system that is an intersection of M matroids and showed that greedy is an $(M+1)$-approximation algorithm. This result yields a 2-approximation factor when $(E, \mathcal{F})$ is a matroid. The authors also considered a simpler variant of the greedy algorithm, that they refer to as the *locally greedy heuristic*, and showed that this algorithm is also a factor-2 approximation algorithm for the problem (P) when $(E, \mathcal{F})$ is a partition matroid. Subsequently, Conforti and Cornuéjols [19] studied the problem (P) over a matroid, but for a richer class of objective functions, f, by introducing the notion of *total curvature* to characterize a set function. They showed that the performance of the greedy algorithm for maximizing a nondecreasing submodular set function of total curvature α is an $(\alpha + 1)$-approximation. Moreover, by showing that $0 \leq \alpha \leq 1$ for nondecreasing submodular functions and $\alpha = 0$ if and only if the function is linear, they generalized the results of Rado-Edmonds and Fisher et al. [38] regarding the performance of the greedy algorithm.

Wolsey [107] considered the problem (P) over an independence system $(E, \mathcal{F})$ given by:

$$\mathcal{F} = \{S \subseteq E : \sum_{e \in S} w_e \leq W\}$$

where w_e, for each $e \in E$, are nonnegative weights and W is a nonnegative integer. This system is simply the set of all feasible solutions to a knapsack constraint, and exemplifies independence systems where $\mathcal{F}$ may be exponentially large, and yet may be encoded succinctly in a problem instance. In what follows, we will see examples where the ground set E itself may be exponentially large and yet may be encoded concisely in a problem instance. Extending a result of Khuller et al. [64] regarding the performance of a *greedy with partial enumeration* algorithm for the BUDGETED MAXIMUM COVERAGE problem, Sviridenko [99] showed that this algorithm is also an $(e/(e-1))$-approximation algorithm for the problem (P) over a knapsack independence system. The $(e/(e-1))$-approximation results of Sviri-

denko [99] and Nemhauser et al. [80] for their respective problems are in fact *best possible* for any polynomial-time approach, unless P=NP (Feige [35]).

Upon the completion of this work, we learnt that recently Calinescu et al. [14] have developed a pipage rounding based $(e/(e-1))$-approximation algorithm for the case of problem (P) where $(E, \mathcal{F})$ is a matroid and f is a sum of weighted rank functions of matroids, which are a rich subclass of monotone submodular functions. Moreover, the authors also give a somewhat different proof for the performance of the standard greedy algorithm with an approximate oracle for the problem (P) when $(E, \mathcal{F})$ is a p-independent family.

2.2.   Motivation

Whereas all of the works highlighted in the previous section, with the exception of the recent paper of Calinescu et al. [14], assume the existence of a polynomial-time procedure (or incremental oracle) in the greedy algorithm to find an optimal incremental element in each iteration, such an oracle may not always be available. We motivate this scenario via an example where the ground set, E, itself may be exponentially large. Consider the following problem studied by Fleischer et al. [39]:

SEPARABLE ASSIGNMENT [39]

Instance: A set, U, of n items and a set, B, of m bins. Each bin $i \in B$ has an independence system $\mathcal{I}_i$ of subsets of items that fit in bin i. A profit p_{ij} for assigning item j to bin i.

Task: Find a subset of items, $S \subseteq U$, and an assignment of these items, $S_i \in \mathcal{I}_i$ to bin i, $S_i \cap S_k = \emptyset$ for $i \neq k$, so as to maximize profit, $\sum_{i \in B} \sum_{j \in S_i} p_{ij}$.

Observe that the family of feasible subsets for each bin i, $\mathcal{I}_i$, is an independence system. Also note that therefore, the constraints defining feasible packings for bin i, implicit in $\mathcal{I}_i$, are *separable* from the constraints for bin j, i.e., the set of feasible packings of bin i are unaffected by the set of feasible packings of bin j. Finally, the authors assume the existence of an α-approximation algorithm for the single-bin subproblem for each bin i: select a feasible packing of items from $\mathcal{I}_i$ of maximum profit. As an example, the GENERALIZED ASSIGNMENT

problem is a special case of the SEPARABLE ASSIGNMENT problem, where the single bin subproblem is the KNAPSACK problem. Specifically, in the GENERALIZED ASSIGNMENT problem, items also have sizes w_{ij} corresponding to each bin i, and each bin itself is a knapsack of a particular capacity B_i. Hence, the single-bin subproblem corresponding to bin i for GENERALIZED ASSIGNMENT would be to find a maximum profit subset of items that fits in bin i. However, for other special cases of SEPARABLE ASSIGNMENT, the single-bin subproblem may be characterized by other forms of resource packing problems, such as the RECTANGLE PACKING or the 2-DIMENSIONAL KNAPSACK problem.

As also noted independently by Chekuri and by Fleischer et al. [39], this problem is an instance of maximizing a normalized, nondecreasing, submodular function over a (partition) matroid. For the sake of completeness, we describe this transformation here.

Observation 2.1. SEPARABLE ASSIGNMENT *is an instance of maximizing a monotone submodular function over a partition matroid.*

Proof. For any instance of the SEPARABLE ASSIGNMENT problem, define a ground set $E = \cup_{i \in B} E_i$, with an element $e_S \in E_i$ corresponding to each feasible packing, $S \in \mathcal{I}_i$, of bin i. The constraints on SEPARABLE ASSIGNMENT now transform to picking at most one element from each set E_i. Let $F = \cup_{i \in B} F_i$, where $F_i \subseteq E_i$ and $|F_i| \leq 1$, represent the set of elements picked. This underlying constraint structure is therefore a partition matroid. However, note that the packings of bins corresponding to any set F, may contain multiple copies of the same item. Therefore it is important that one does not double-count the profit for these items. This may be taken care of by writing the objective function as:

$$f(F) = \sum_{j \in U} \max\{p_{ij} : i \in B, e_S \in F_i, j \in S\} \ .$$

Observe that this definition of f extends to all subsets $F \subseteq E$, even when $|F_i|$ is more than 1. Observe also that the summation is over all items $j \in U$, and the maximum is over all bins i that contain item j. In other words, if an item j is in multiple bins corresponding to a set, F, then out of all the bins, i, that item j is in, we assign only the maximum p_{ij}

value to item j. It is not hard to verify that indeed this function f is nondecreasing, and has decreasing marginal values. Suppose that $S \in \mathcal{I}_i$ is a feasible packing for bin i. Observe that the incremental value of element e_S to a set F, $\rho_{e_S}(F)$ is given by:

$$\rho_{e_S}(F) = f(F \cup \{e_S\}) - f(F) = \sum_{j \in S} \max\{p_{ij} - \max\{p_{kj} : e_P \in F, P \in \mathcal{I}_k, j \in P\}, 0\}.$$

Intuitively, the incremental value of an element is the incremental profit value of the set of items in the corresponding packing that the element represents. As the set F grows, the likelihood of the items in the packing, S, of bin i having the highest profit, p_{ij}, decreases and hence, f has decreasing marginal values. Therefore, f is submodular. ∎

We have seen that in the underlying matroid, an element of the ground set corresponds to a feasible packing in a bin. Consequently, the role of an incremental oracle in the greedy algorithm for this problem is to pick a feasible packing among all feasible packings of maximum incremental value to the existing solution. This would typically involve solving a knapsack problem (or even a rectangle packing problem), since the set of all feasible packings might be exponentially large. However, since such packing problems are typically NP-hard, we cannot hope to have an optimal incremental oracle, unless P=NP. Hence, generalized results such as the one described below, assuming instead the existence of an α-approximation oracle to find a "good" incremental element, are in order.

In this chapter, we present bounds on the performance of a greedy algorithm that uses an α-approximation algorithm as the incremental oracle to determine an incremental element to add to the greedy solution. We summarize our main results in the following section.

2.3. Our Results

We begin by considering the problem of maximizing a nondecreasing submodular function over uniform matroids. Generalizing a previous result due to Nemhauser et al. [80], we show that in the presence of an α-approximate incremental oracle, the standard greedy algorithm

55

is an $\left(\frac{e^{\frac{1}{\alpha}}}{e^{\frac{1}{\alpha}}-1}\right)$-approximation algorithm for this problem. Further, we also discuss how our result generalizes similar previous results due to Hochbaum and Pathria [51] in the context of the MAXIMUM COVERAGE problem, and Chekuri and Khanna [15] with regards to the MULTIPLE KNAPSACK problem with identical bin capacities.

Partition matroids generalize uniform matroids, in that the ground set E contains elements of different kinds, with individual restrictions on how many elements may be selected of each kind. In Section 2.5, we consider a variant of the standard greedy algorithm, namely the locally greedy algorithm, previously proposed by Fisher et al. [38] and consider the performance of this algorithm for maximizing nondecreasing submodular set functions over partition matroids. Extending a result of Fisher et al. [38] to α-approximate incremental oracles, we show that the locally greedy algorithm guarantees a tight factor-$(\alpha + 1)$ result for the submodular function maximization problem over partition matroids. We also show that various optimization problems that arise in the context of the winner determination in combinatorial auctions (Lehmann et al. [71], Dobzinski and Schapira [31]), generalized assignment problems (Fleischer et al. [39]), basic utility games (Vetta [101], Mirrokni and Vetta [75]), wireless networks (Abrams et al. [2]), etc. may be cast into the framework of maximizing a submodular function over a partition matroid. Consequently, we reinterpret and unify the results pertaining to these problems within our framework.

Adapting a randomized algorithm proposed by Dobzinski and Schapira [31] for the winner determination problem in combinatorial auctions with submodular bidders, we propose a randomized algorithm for maximizing a submodular function over a partition matroid, in Section 2.6. We show that in expectation, this algorithm guarantees a $(2 - \frac{1}{n})$-approximate solution for the problem, where $n = \max_i |E_i|$. If the size of the ground set, $|E|$, is polynomial in the size of the input, we show that this algorithm also runs in polynomial time. Finally, we show that this algorithm implies a new polynomial-time randomized $\frac{3}{2}$-approximation algorithm for MAX SAT, interestingly matching the performance of Johnson's algorithm for the problem (Johnson [60], Chen et al. [17]). Indeed, there do exist improved linear and semidefinite programming-based approximation algorithms for MAX SAT, but our approach

serves to illustrate the benefits of a more generalized study of optimization problems.

In Section 2.7, we consider the problem of maximizing a submodular function over an independence system. If the independence system is an intersection of a finite number, M, of matroids, then Fisher et al. [38] show that the greedy algorithm with an optimal incremental oracle is an $(M + 1)$-approximation algorithm for this problem. When only an α-approximate incremental oracle is available, we show that the greedy algorithm is an $(\alpha M + 1)$-approximation for the problem. Based on this result, we improve upon a previous result of Fleischer et al. [39] for the k-MEDIAN WITH HARD CAPACITIES AND PACKING problem and we present a greedy $(\alpha+1)$-approximation algorithm for it. Finally, we conclude in Section 2.8 by highlighting some interesting open questions and future directions that result from this work.

2.4. Generalized Results over Uniform Matroids

Let $(E, \mathcal{F})$ be a uniform matroid, i.e., $\mathcal{F} = \{S \subseteq E : |S| \leq k\}$ for some integer k. Consider the problem of maximizing a normalized, nondecreasing, submodular function f over this uniform matroid. Using notation introduced by Farahat and Barnhart [34], we represent this problem as $f_S|\mathcal{F}_U$. We describe a generalized greedy algorithm below that uses an α-approximation algorithm as the incremental oracle to find an element e with the best incremental value, $\rho_e(S) = f(S \cup \{e\}) - f(S)$. Note that in the case of uniform matroids, the role of the admissibility oracle in the greedy algorithm is trivial – as long as the size of the solution set, S, is strictly smaller than k, any element is admissible.

GREEDY ALGORITHM FOR $f_S|\mathcal{F}_U$

Step 1: Set $i = 1$; let $S_0 = \emptyset$.

Step 2: Select an element $e_i \in E$ for which $\alpha \cdot \rho_{e_i}(S_{i-1}) \geq \max_{e \in E \setminus S_{i-1}} \rho_e(S_{i-1})$ using an α-approximate incremental oracle.

Step 3: Set $S_i = S_{i-1} \cup \{e_i\}$.

Step 4: Set $i := i + 1$. If $i \leq k$, then goto **Step 2**.

S_i represents the set generated by the greedy algorithm after i iterations. Let $S^G = S_k$ be the solution returned by the greedy algorithm. Let ρ_i represent the incremental profit obtained by the addition of element e_i to the set S_{i-1}. Let ρ'_i represent the optimal incremental profit that could have been obtained, given the set S_{i-1} was selected by the first $i-1$ iterations of the greedy algorithm. Since we use an α-approximation oracle in order to determine the element with the best incremental objective function value, it follows that $\rho_i \leq \rho'_i \leq \alpha \cdot \rho_i$. Observe that if one had access to an optimal incremental oracle, then it would be the case that $\rho_i \geq \rho_{i+1}$. However, since we only use an approximate incremental oracle, this need not hold anymore. Thus, *the use of an approximate incremental oracle does not preserve the nonincreasing property of incremental values of elements selected by the greedy algorithm.* This is a recurring theme throughout the paper and adds a measure of complexity to the original proofs of Nemhauser et al. [80] and Fisher et al. [38].

We begin with the following characterization for nondecreasing submodular functions, and present its proof for the sake of completeness:

Lemma 2.2 (Nemhauser et al. [80]). *f is a nondecreasing submodular set function on E if and only if $f(T) \leq f(S) + \sum_{j \in T \setminus S} \rho_j(S)$ for all $S, T \subseteq E$.*

Proof. Suppose that f is a nondecreasing, submodular set function defined on E. Furthermore, for any two subsets, $S, T \subseteq E$, let $T \setminus S = \{e_1, e_2, \ldots, e_k\}$. In other words, $S \cup T = S \cup \{e_1, e_2, \ldots, e_k\}$. Now, since f is nondecreasing, we have that:

$$
\begin{aligned}
f(T) \quad &\leq \quad f(S \cup T) \\
&= \quad f(S) + (f(S \cup \{e_1\}) - f(S)) + (f(S \cup \{e_1, e_2\}) - f(S \cup \{e_1\})) + \ldots \\
&\quad + (f(S \cup \{e_1, e_2, \ldots, e_k\}) - f(S \cup \{e_1, \ldots, e_{k-1}\})) \\
&= \quad f(S) + \rho_{e_1}(S) + \rho_{e_2}(S \cup \{e_1\}) + \ldots + \rho_{e_k}(S \cup T \setminus \{e_k\}) \\
&\leq \quad f(S) + \rho_{e_1}(S) + \rho_{e_2}(S) + \ldots + \rho_{e_k}(S) \qquad\qquad (2.2) \\
&= \quad f(S) + \sum_{e \in T \setminus S} \rho_e(S).
\end{aligned}
$$

58

Inequality (2.2) follows from the decreasing marginal values definition of submodularity, (**F2a**).

Conversely, suppose that a function f satisfies the inequality, $f(T) \leq f(S) + \sum_{e \in T \setminus S} \rho_e(S)$ for all $S, T \subseteq E$. Suppose that $A \subseteq B \subseteq E$ and $e \in E \setminus B$. In addition, suppose that $B \setminus A = \{e_1, e_2, \ldots, e_k\}$. Now, substituting, $T = A, S = B$, we have that $f(A) \leq f(B)$, since $A \setminus B = \emptyset$, and therefore f is nondecreasing. Moreover, substituting $T = A \cup \{e_1, e\}$ and $S = A$, we have that:

$$f(A \cup \{e_1, e\}) \leq f(A) + \rho_e(A) + \rho_{e_1}(A) = f(A \cup \{e_1\}) + \rho_e(A),$$

which in turn implies that:

$$f(A \cup \{e_1, e\}) - f(A \cup \{e_1\}) \leq \rho_e(A) \Rightarrow \rho_e(A \cup \{e_1\}) \leq \rho_e(A).$$

Using one-element increments, one may similarly show that $\rho_e(A \cup \{e_1, e_2\}) \leq \rho_e(A \cup \{e_1\})$, and so on. Putting these together, we would therefore have that $\rho_e(B) \leq \rho_e(A \cup \{e_1, e_2, \ldots, e_{k-1}\}) \leq \ldots \leq \rho_e(A \cup \{e_1\}) \leq \rho_e(A)$. Hence it follows that f is nondecreasing and submodular. ∎

Suppose that $z_{opt} = \max_{S \subseteq E} \{f(S) : |S| \leq k\}$, with f is normalized, nondecreasing, and submodular. We then show that:

Theorem 2.3. *If z_g is the value of the* GREEDY ALGORITHM FOR $f_S | \mathcal{F}_U$*, then* $\frac{z_{opt}}{z_g} \leq \frac{(\alpha k)^k}{(\alpha k)^k - (\alpha k - 1)^k} \leq \frac{e^{\frac{1}{\alpha}}}{e^{\frac{1}{\alpha}} - 1}$.

Proof. Suppose that S^G is the set generated by the greedy algorithm and T is an optimal solution to the above problem. Let ρ'_i represent the best incremental value that could have been obtained during the i^{th} iteration of the greedy algorithm. By substituting $S = \emptyset$ in Lemma 2.2 and observing that $|T| \leq k$, it follows that:

$$z_{opt} = f(T) \leq \sum_{j \in T} f(j) \leq k\rho'_1 \leq k(\alpha \rho_1).$$

59

Now, applying Lemma 2.2 to the solution of the greedy algorithm after j iterations, S_j, implies that:

$$z_{opt} \leq f(S_j) + \sum_{i \in T \setminus S_j} \rho_i(S_j).$$ (2.3)

Given that $f(S_j) = \sum_{i=1}^{j} \rho_i$ and that

$$\alpha \rho_{j+1} \geq \rho'_{j+1} \geq \rho_i(S_j) \quad \text{for all } i \in E \setminus S_j,$$

equation (2.3) now yields the following inequality:

$$z_{opt} \leq \sum_{i=1}^{j} \rho_i + k \cdot (\alpha \rho_{j+1}),$$

which implies that

$$\rho_{j+1} \geq \frac{1}{\alpha k} z_{opt} - \frac{1}{\alpha k} \sum_{i=1}^{j} \rho_i.$$

Adding $\sum_{i=1}^{j} \rho_i$ on both sides of the above inequality, we get an inequality of the form:

$$\sum_{i=1}^{j+1} \rho_i \geq \frac{1}{\alpha k} z_{opt} + \frac{(\alpha k - 1)}{\alpha k} \sum_{i=1}^{j} \rho_i.$$ (2.4)

We now prove by induction on j that:

$$\sum_{i=1}^{j} \rho_i \geq \frac{(\alpha k)^j - (\alpha k - 1)^j}{(\alpha k)^j} \cdot z_{opt}.$$

For $j = 1$, we have that $\rho_1 \geq \frac{1}{\alpha k} \cdot z_{opt}$. Assume that the claim holds for $j - 1$. Now, applying the induction hypothesis on equation (2.4), we have:

$$\sum_{i=1}^{j} \rho_i \geq \frac{1}{\alpha k} \cdot z_{opt} + \frac{\alpha k - 1}{\alpha k} \cdot \frac{(\alpha k)^{j-1} - (\alpha k - 1)^{j-1}}{(\alpha k)^{j-1}} \cdot z_{opt}.$$ (2.5)

Simplifying the right-hand side of the above expression yields the induction claim. Finally,

setting $j = k$ we have:

$$z_g = \sum_{i=1}^{k} \rho_i \geq \frac{(\alpha k)^k - (\alpha k - 1)^k}{(\alpha k)^k} \cdot z_{opt},$$

which proves the approximation ratio claim that :

$$\frac{z_{opt}}{z_g} \leq \frac{(\alpha k)^k}{(\alpha k)^k - (\alpha k - 1)^k} \leq \frac{e^{\frac{1}{\alpha}}}{e^{\frac{1}{\alpha}} - 1}.$$

■

The above result essentially follows in a manner similar to that of Nemhauser et al. [80] and serves to point out the effect of α on the approximation factor of the greedy algorithm. For the case when $\alpha = 1$, the result is precisely that of Nemhauser et al. [80] and therefore tight. Theorem 2.3 also generalizes a similar result due to Hochbaum and Pathria [51] (see also Hochbaum [50]) in the context of the MAXIMUM COVERAGE problem and its applications. We discuss this in greater detail in subsequent sections.

2.4.1 Discussion on Running Time of Greedy Algorithm

Denote the running time of the α-approximate incremental oracle by P. It follows then that the running time of the GREEDY ALGORITHM FOR $f_S|\mathcal{F}_U$ is $O(k\text{P})$, where at most k elements need to be selected from the uniform matroid. Observe that the running time of the algorithm itself does not depend on the size of the ground set, E, which may possibly be exponentially large. As discussed earlier, and motivated in SEPARABLE ASSIGNMENT, in certain underlying problems, the ground set E may be encoded concisely, even though it is exponentially large. However, without loss of generality, any constant-factor approximation algorithm to $f_S|\mathcal{F}_U$ must select $O(k)$ elements (else consider a modular objective function with equal weights for all elements). Hence in that sense, the greedy algorithm is an efficient algorithm as long as P is polynomial in the input size of the underlying problem. Moreover, for problems such as the relaxed AD PLACEMENT problem (Section 1.7) where the number

of elements to be selected is itself encoded using $\log k$ bits, it is often possible to modify the greedy algorithm appropriately so that its running time is still polynomial, as we had showed in Section 1.7.1 for the relaxed AD PLACEMENT problem.

2.4.2 Applications of Generalized Result

We begin by studying implications of Theorem 2.3 for the MAXIMUM COVERAGE problem in discrete optimization. The MAXIMUM COVERAGE problem may be stated as follows:

MAXIMUM COVERAGE

Instance: A set of elements, U, a collection $\mathcal{R}$ of subsets of U, and an integer k. A nonnegative profit, p_j, corresponding to each element $j \in U$.

Task: Select k subsets $U_1, \ldots, U_k$ of U, with each $U_i \in \mathcal{R}$, such that the profit of the elements in $\cup_{i=1}^{k} U_i$ is maximized.

Vohra and Hall [105] noted that MAXIMUM COVERAGE is indeed a special case of $f_S | \mathcal{F}_U$. For completeness, we present this transformation below.

Observation 2.4. MAXIMUM COVERAGE *is a special case of* $f_S | \mathcal{F}_U$.

Proof. Consider an underlying uniform matroid $(E, \mathcal{F})$ where each element $i \in E$ corresponds to a subset $U_i \in \mathcal{R}$, and $\mathcal{F}$ is the collection of all subsets of E of size at most k. This matroid would characterize the underlying constraint in MAXIMUM COVERAGE. Moreover, the objective of MAXIMUM COVERAGE may be rewritten as the following function f:

$$f(F) = p(\cup_{i \in F} U_i), \quad \text{where } p(S) = \sum_{j \in S} p_j \text{ for each } S \subseteq U.$$

It follows from this definition that f is normalized, nondecreasing with decreasing marginal values, and therefore submodular. ∎

Hochbaum and Pathria [51] present a greedy algorithm to solve MAXIMUM COVERAGE, and a scenario where finding a subset that gives maximum improvement might be hard. They obtain the same bound as the one in Theorem 2.3, assuming that one is able to pick an

α-approximate solution in each stage. Hochbaum and Pathria [51] also describe a number of applications that can be modeled as the MAXIMUM COVERAGE problem in a setting of approximate improvement.

Theorem 2.3 also implies the bound of the performance of the greedy algorithm obtained by Chekuri and Khanna [15] for the MULTIPLE KNAPSACK problem with identical bin capacities. The MULTIPLE KNAPSACK problem may be stated as follows:

MULTIPLE KNAPSACK

Instance: Nonnegative integers, n, m, $p_1, \ldots, p_n$, $w_1, \ldots, w_n$, and $W_1, \ldots, W_m$.

Task: Find m subsets $S_1, \ldots, S_m \subseteq \{1, \ldots, n\}$, $S_i \cap S_k = \emptyset$ for $i \neq k$, such that $\sum_{j \in S_i} w_j \leq W_i$ for $i = 1, \ldots, m$ and $\sum_{i=1}^{m} \sum_{j \in S_i} p_j$ is maximum.

In the case that all m bins have the same capacity, $W_1 = W_2 = \ldots = W_m = W$, this problem is an instance of $f_S|\mathcal{F}_U$. The transformation for this is essentially identical to that described for SEPARABLE ASSIGNMENT in Observation 2.1. Using an FPTAS for the KNAPSACK problem as an incremental oracle, the greedy $(\frac{e}{e-1} + \epsilon)$-approximation result of Chekuri and Khanna [15] follows from Theorem 2.3.

2.5. The Locally Greedy Algorithm and Partition Matroids

In this section, we generalize the performance bounds of a special version of the greedy algorithm, namely the *locally greedy heuristic* of Fisher et al. [38], to maximize a submodular function over a partition matroid. Recall that a partition matroid, $(E, \mathcal{F})$ is given by $\mathcal{F} = \{F : F = \cup_{i=1}^{k} F_i$ where $F_i \subseteq E_i, |F_i| \leq l_i$ for $i = 1, \ldots, k\}$. We assume that we only have at our disposal an α-approximation algorithm to play the role of an incremental oracle for each element type, E_i. We shall refer to this problem as $f_S|\mathcal{F}_P$, where the subscript P denotes the partition matroid. The locally greedy algorithm for this problem is as follows:

LOCALLY GREEDY ALGORITHM FOR $f_S|\mathcal{F}_P$

Step 1: Set $i := 1$; let $S_0 := \emptyset$; $m := 1$.

Step 2: Set $j := 1$.

Step 3: Select an element $e_j \in E_i$ for which $\alpha \cdot \rho_{e_j}(S_{m-1}) \geq \max_{e \in E_i \setminus S_{m-1}} \rho_e(S_{m-1})$ using an α-approximation algorithm as the incremental oracle of type i.

Step 4: Set $S_m := S_{m-1} \cup \{e_j\}$; $m := m + 1$.

Step 5: Set $j := j + 1$. If $j \leq l_i$, then goto **Step 3**.

Step 6: Set $i := i + 1$. If $i \leq k$, then goto **Step 2**.

In the above algorithm, i is a counter of the type, E_i, of elements in consideration; j is a counter for the number of elements selected within a particular type; and m represents the number of elements selected by the greedy algorithm at any point. The locally greedy algorithm basically begins by selecting "profitable" elements of the first type, E_1, from E, until it has picked l_1 elements from E_1, and then it proceeds to do so for the second type of elements, E_2 in E and so on. Thus, the number of elements, m, in the greedy solution, S^G, is at most $\sum_{i=1}^{k} l_i$.

It is important to note that, *the order in which the locally greedy algorithm deals with elements of different types is completely arbitrary.* Furthermore, the incremental oracle in the locally greedy algorithm only need select *an approximate best element within each particular type*, rather than an approximate best element across all element types. Of course, with an α-approximate incremental oracle over each type, one may simulate an α-approximate incremental oracle over all types in $O(k)$ time, by taking the α-best element of each of k types and selecting the best of them.

We now present our main result for partition matroids.

Theorem 2.5. *If z_g is the value of the solution provided by the* LOCALLY GREEDY ALGORITHM FOR $f_S|\mathcal{F}_P$, *and z_{opt} is the value of an optimal solution, then* $\frac{z_{opt}}{z_g} \leq \alpha + 1$.

Proof. Let S^G represent the greedy solution and T an optimal solution to $f_S|\mathcal{F}_P$. Substitut-

ing them in Lemma 2.2, we have that:

$$z_{opt} = f(T) \leq f(S^G) + \sum_{j \in T \setminus S^G} \rho_j(S^G).$$

Now, suppose $T \setminus S^G = \cup_{i=1}^k T_i$, where $T_i \subseteq E_i$. Also, suppose $S^G = \cup_{i=1}^k S_i^G$ where $S_i^G \subseteq E_i$. Let e_i be the element in S_i^G that was selected with the lowest ρ value, which was at a point in the algorithm when the current greedy solution, just before the addition of e_i, was S^{e_i}. Mathematically,

$$\rho_{e_i}(S^{e_i}) = \min_{e \in S_i^G} \rho_e(S^e) \ .$$

In other words, $\rho_{e_i}(S^{e_i})$ is the minimum incremental value selected by the greedy algorithm among elements in S_i^G. But given that we are using an α-approximation algorithm, we also know that:

$$\alpha \cdot \rho_{e_i}(S^{e_i}) \geq \rho_e(S^{e_i}) \quad \text{for all } e \in E_i \setminus S^{e_i}. \tag{2.6}$$

Also, since $\sum_{j \in T \setminus S^G} \rho_j(S^G) = \sum_{i=1}^k \sum_{j \in T_i} \rho_j(S^G)$, it now follows that:

$$
\begin{aligned}
z_{opt} &\leq f(S^G) + \sum_{i=1}^k \sum_{j \in T_i} \rho_j(S^G) \\
&\leq z_g + \sum_{i=1}^k \sum_{j \in T_i} \rho_j(S^{e_i}) \ .
\end{aligned}
\tag{2.7}
$$

Inequality (2.7) follows from the submodularity of f (see property (**F2a**)), since $S^{e_i} \subseteq S^G$, for $i = 1, \ldots, k$. Noting that $T_i \subseteq E_i \setminus S^G \subseteq E_i \setminus S^{e_i}$ and using inequality (2.6), the right-hand side of inequality (2.7) yields further that:

$$
\begin{aligned}
z_{opt} &\leq z_g + \sum_{i=1}^k \sum_{j \in T_i} \alpha \cdot \rho_{e_i}(S^{e_i}) = z_g + \sum_{i=1}^k \alpha \cdot |T_i| \cdot \rho_{e_i}(S^{e_i}) \\
&\leq z_g + \alpha \cdot z_g \\
&\leq (\alpha + 1) \cdot z_g
\end{aligned}
\tag{2.8}
$$

Inequality (2.8) is implied from the way we picked e_i to be the element with the lowest incremental function value in S_i^G, and since we may assume without loss of generality that $|T_i| \leq |S_i^G| \leq l_i$ (recall that $(E, \mathcal{F})$ is a matroid and f is nondecreasing, and therefore we can always add elements to S_i^G so that $|T_i| \leq |S_i^G|$). ∎

2.5.1 Discussion on Running Time of Locally Greedy Algorithm

It is easy to see that the running time of the LOCALLY GREEDY ALGORITHM FOR $f_S|\mathcal{F}_P$ only depends on the number of elements, l_i, to be picked of each type i, and the running time of the incremental oracle of type i, P_i, as $O(\sum_{i=1}^{k} l_i \mathsf{P}_i)$. It does not depend on the size of the ground set, E, which may potentially be exponentially large, as in SEPARABLE ASSIGNMENT. Moreover, as noted in the running time of the GREEDY ALGORITHM FOR $f_S|\mathcal{F}_U$, since any constant-factor approximation algorithm selects a solution of size $O(\sum_{i=1}^{k} l_k)$, it follows that the locally greedy algorithm is an efficient algorithm as long as the running time, P_i, of each oracle is polynomial in the size of the input of the underlying optimization problem.

If one uses an optimal incremental oracle in the greedy algorithm, implying $\alpha = 1$, Theorem 2.5 matches the result of Fisher et al. [38], and guarantees a bound of 2 for the performance of a locally greedy algorithm over $f_S|\mathcal{F}_P$. This result of Fisher et al. [38] for the locally greedy heuristic seems relatively unknown compared to their result for the performance of the standard greedy algorithm for $f_S|\mathcal{F}_M$, where the subscript M denotes arbitrary matroids. In the following subsections, we reinterpret some results in recent literature based on the result of Fisher et al. [38] for the locally greedy heuristic. But first, we begin by illustrating an application of the theorem to MAX SAT.

2.5.2 Max Sat

The SATISFIABILITY problem is one the most seminal problems in theoretical computer science, and was the first problem shown to be NP-complete (Cook [20]). The optimization version of the problem, MAXimum SATisfiability, may be stated as follows:

> MAX SAT
>
> *Instance*: A set X of n Boolean variables, a family $\mathcal{Z}$ of clauses, each a disjunction of literals over X, and a weight function $c : \mathcal{Z} \to \mathbb{R}_{\geq 0}$.
>
> *Task*: Find a truth assignment T of X such that the total weight of the clauses in $\mathcal{Z}$ that are satisfied by T is maximum.

We first show that MAX SAT is a special case of $f_S | \mathcal{F}_P$. Corresponding to any instance of MAX SAT, define an independence system, $(E, \mathcal{F})$, where the ground set, $E = \cup_{i=1}^n E_i$, contains elements of n different types. Each type of elements, E_i, represents all possible truth assignments of a Boolean variable, $x_i \in X$. More specifically:

$$E_i = \{e_i^T, e_i^F\} \text{ for each } x_i \in X,$$

where e_i^T (e_i^F, respectively) would correspond to x_i being assigned the value *true* (*false*, respectively). In any feasible truth assignment to MAX SAT, each variable, x_i, may be assigned at most one truth value (*true* or *false*). This constraint defines a partition matroid over E, wherein at most one element may be picked of each type, E_i.

We claim that the objective of maximizing the weight of satisfied clauses may be rewritten as the following function, f, defined over subsets of E:

$$f(S) = \sum_{z \in \mathcal{Z}_S} c(z) \text{ where } \mathcal{Z}_S = \{z \in \mathcal{Z} : x_i \in z, e_i^T \in S \text{ or } \bar{x}_i \in z, e_i^F \in S\}$$

This may be easily verified by considering the set S corresponding to any given feasible truth assignment. Furthermore, f is normalized and nondecreasing, since $\mathcal{Z}_S$ is nondecreasing in S, and since $c(z) \geq 0$. Moreover, without loss of generality, the incremental value of any element $e_i^T \in E$ to any set S would correspond to the weight of all additional clauses containing the literal x_i that were not already satisfied by literals corresponding to the elements by S. From this interpretation, it follows that f has the property of decreasing marginal values, and is therefore submodular. Finally, note that since f is nondecreasing, corresponding to any solution to $f_S | \mathcal{F}_P$, one may construct a new solution that is a basis with at least the

same objective value as the original solution. Moreover, any basis solution corresponds to a feasible truth assignment, thus completing the claim.

It is not hard to see that the locally greedy algorithm for MAX SAT would then correspond to arbitrarily deciding on an ordering of the variables, and iteratively assigning the truth value, *true* or *false*, to each variable x_i that has maximum incremental value (sum of weight of additional clauses satisfied) to the solution at hand. The result of Fisher et al. [38] would imply that this algorithm has an approximation guarantee of 2 for MAX SAT. Of course, the locally greedy algorithm does not exploit the complete underlying structure of the objective function, aside from submodularity. Consequently, there are a number of improved approximation algorithms for this problem, based on linear programming, semidefinite programs, and the probabilistic method. We refer the reader to standard texts in approximation algorithms such as [67, 100] for further exposition on MAX SAT.

2.5.3 Winner Determination in Combinatorial Auctions

Combinatorial auctions are mechanisms via which multiple non-identical items are sold to bidders who express preferences over combinations of items, and not just single items. Such auctions assume particular relevance when the items being sold are either *complements* or *substitutes* to each other. In particular, Lehmann et al. [71] study the problem of an auctioneer who would like to allocate a set of items, X, of decreasing marginal values amongst n *submodular* bidders so as to maximize total social welfare. More formally, the problem may be stated as follows:

WINNER DETERMINATION [71]

Instance: A set X of items; n bidders, each bidder j having a submodular valuation function, $v_j : 2^X \to \mathbb{R}_{\geq 0}$ which is normalized and nondecreasing.

Task: A partition of the items X, into pairwise disjoint sets, $S_1, \ldots, S_n$, so as to maximize $\sum_{j=1}^n v_j(S_j)$.

It has been observed in the literature that the WINNER DETERMINATION problem is a special case of $f_S | \mathcal{F}_P$. We present the transformation below.

Consider a ground set given by $E = \cup_{i \in X} E_i$, where each element $e_{ij} \in E_i$ corresponds to allocating item i to bidder j. The constraint defined by the WINNER DETERMINATION problem, that any item $i \in X$ may be assigned to at most one bidder, would therefore transform to picking at most element from each set E_i. Clearly, the set of all feasible subsets of E defined by this constraint would therefore be a partition matroid. Moreover, the objective function of the auctioneer is to maximize social utility, $\sum_{j=1}^{n} v_j(S_j)$ where $(S_1, \ldots, S_n)$ is a partition of X. This objective is also nondecreasing and submodular, since it is a sum of nondecreasing submodular functions. Based on any set $F \subseteq E$, define

$$S_j^F = \{i | e_{ij} \in F\}.$$

Furthermore, the objective function of the auctioneer may be rewritten as $f(F) = \sum_{j=1}^{n} v_j(S_j^F)$, which is clearly monotone submodular on the base set, E. This may be easily verified by noting that the marginal value of any element e_{ij} to a set F is indeed nonincreasing, since the marginal value of the corresponding item i being allocated to bidder j is itself nonincreasing. Thus the WINNER DETERMINATION problem is indeed an instance of $f_S | \mathcal{F}_P$, as also noted by Lehmann et al. [71].

Interestingly, we point out that the factor-2 greedy approximation algorithm proposed by Lehmann et al. [71] for this problem turns out to be exactly the locally greedy algorithm of Fisher et al. [38], as both algorithms are independent of item ordering. To complete the analogy between both these greedy algorithms, we observe that Lehmann et al. [71] assume access to a *value oracle* for each player, to encode the submodular valuations of the players. That is, given a set S of items, the value oracle for bidder j outputs $v_j(S)$. In the corresponding locally greedy algorithm, the existence of an optimal incremental oracle follows from the existence of the value oracle, and from the fact that there are only n bidders whose valuations need to be checked to find the best incremental element. The authors do not make this connection and instead claim that the family of greedy algorithms that they consider is wider than that of Fisher et al. [38].

One might observe that in the transformation described earlier in this section, the structure of the submodular objective function is "separable," in the sense that the marginal utility of any item e_{ij} only depends on which other elements have been allocated to bidder j, i.e., elements of the form e_{kj} in the current solution, $F \subseteq E$. It is conceivable that one might leverage this special structure to devise improved approximation algorithms for the WINNER DETERMINATION problem. Indeed, Dobzinski and Schapira [31] propose a $(2 - \frac{1}{n})$-approximation algorithm for this problem, where n represents the number of elements of each type. However, we show in Section 2.6, that their algorithm may be adapted to provide a $(2 - \frac{1}{n})$-approximation algorithm for $f_S|\mathcal{F}_P$ as well. Interestingly, Khot et al. [63] recently showed that there is no polynomial time approximation algorithm with a factor better than $e/(e - 1)$ for the WINNER DETERMINATION problem, unless P=NP.

2.5.4   Convergence Issues in Competitive Games

Vetta [101] studied the following strategic game played amongst n players: associated with each player j is a disjoint ground set V_j of *actions*, and S_j, a collection of subsets of V_j. Any set $s_j \in S_j$ corresponds to a feasible *strategy* of player j. In addition, suppose that $\emptyset_j \in S_j$ corresponds to the null strategy for player j. A strategy profile or *state*, $S = (s_1, \ldots, s_n)$, represents the corresponding strategies being played by each player at a particular time. Let $S \oplus s'_j = (s_1, \ldots, s_{j-1}, s'_j, s_{j+1}, \ldots, s_n)$ denote the state obtained if player j were to change its strategy to s'_j. Assume that $\alpha_j : \Pi_j S_j \to \mathbb{R}$ represents the private utility function of player j, and $\gamma : \Pi_j S_j \to \mathbb{R}$, the social objective function. Suppose that the social objective function, $\gamma(.)$, is a monotone submodular set function defined on $\cup_{j=1}^n V_j$, i.e.,

$$\gamma(S) = g(\cup_{j=1}^n s_j),$$

where g is a monotone submodular function defined on $\cup_j V_j$. Based on different assumptions on γ and α_j, Vetta [101] introduced the following types of games:

- **Utility Game**: A strategic game as described above is said to be a utility game if it

satisfies the *Vickrey condition*:

$$\alpha_j(S) \geq \gamma(S) - \gamma(S \oplus \emptyset_j) \qquad \text{for all feasible states, } S.$$

- **Valid Utility Game**: A valid utility game is a utility game that satisfies the *Cake condition*:

$$\sum_j \alpha_j(S) \leq \gamma(S) \qquad \text{for all feasible states, } S.$$

- **Basic Utility Game**: A basic utility game is a valid utility game that satisfies the Vickrey condition with equality:

$$\alpha_j(S) = \gamma(S) - \gamma(S \oplus \emptyset_j) = g(\cup_i s_i) - g(\cup_{i \neq j} s_i).$$

We reinterpret valid utility games, and in particular, basic utility games, as an equivalent *decentralized approach* to maximizing a submodular function over a partition matroid, where at most one element may be picked of each type. We illustrate this equivalence for basic utility games, and the equivalence for valid utility games follows similarly.

Corresponding to an instance of a basic utility game, construct a ground set $E = \cup E_j$ and add an element e_j in E_j corresponding to each feasible strategy $s_j \in S_j$ of player j. That a player may select at most one strategy from its feasible set of strategies would lead to a natural partition matroid on this underlying ground set, E. Correspondingly, a function, f, may be defined on any subset $F \subseteq E$ as follows:

$$f(F) = g(\cup_{e_j \in F} s_j).$$

It may be verified that f is monotone submodular on E, since g itself is monotone submodular on the underlying ground set $\cup_j V_j$. This would require making use of an alternate characterization of submodular functions, namely that $\rho_A(B) \geq \rho_A(C)$ for all $A \subseteq E$, $B \subseteq C \subseteq E$. Observe that in the transformation so far, we have made no assumptions whatsoever re-

garding the underlying private utilities of the players. Thus, by imbuing each player with any private utility function, we may define a corresponding *partition matroid game*, where E_j would correspond to the strategy space of player j. Indeed, if the private utility of each player in the partition matroid game is set to

$$\alpha_j(F) = f(F) - f(F \setminus E_j) \text{ for any } F \subseteq E,$$

then this partition matroid game defined would in fact be the basic utility game we sought to transform, since this private utility matches the private utility of player j in the basic utility game.

Conversely, starting with any instance of $f_S | \mathcal{F}_P$ wherein at most one element may be picked from each type (we argue later in Section 2.6 that this is without loss of generality), we may similarly define a partition matroid game over the instance. By imbuing the player representing elements of type j with the private utility function,

$$\alpha_j(F) = f(F) - f(F \setminus E_j) \text{ for any } F \subseteq E,$$

we clearly satisfy the Vickrey condition with equality. Moreover, it is not hard to verify that these utilities also satisfy the Cake condition (refer Theorem 2.5 of Vetta [101]). Thus, we may define a basic utility game corresponding to each instance of $f_S | \mathcal{F}_P$ as well. Additionally, by defining alternate private utilities in the partition matroid game, one may draw a similar correspondence to valid utility games. Via this correspondence, we may now reinterpret the results of Vetta [101] for valid utility games as the performance bounds of a decentralized approach to $f_S | \mathcal{F}_P$.

One of the main results of Vetta [101] (Theorem 3.4) is that there exists a Nash equilibrium in any valid utility game, and that the expected social value of any (pure or mixed strategy) Nash equilibrium is at least half the social optimal value. This result may alternately be interpreted as:

Corollary 2.6. *Any Nash equilibrium of a decentralized valid-utility game approach to $f_S | \mathcal{F}_P$*

is a factor 2 approximation to the optimal solution.

Vetta [101] gives examples that imply that this factor 2 result is indeed tight. Unfortunately, Goemans et al. [45] show that for some instances of valid utility games (alternately, in a valid-utility game approach for certain instances of $f_S|\mathcal{F}_P$), finding a Nash (sink) Equilibrium is PLS-complete. However, whether there exist polynomial-time convergence schemes to good equilibrium solutions in alternate decentralized approaches to $f_S|\mathcal{F}_P$ remains an open question.

Interestingly, iterative improved response strategies in a valid-utility game framework for $f_S|\mathcal{F}_P$ closely resemble local search approaches to $f_S|\mathcal{F}_P$. Indeed, Fisher et al. [38] give similar bounds on the performance of an *interchange heuristic*, a local improvement procedure. In an iteration of the interchange heuristic, while there exists an element e outside the current solution, S, that may be swapped with an element in S so as to improve the value of the solution while maintaining independence simultaneously, modify S by interchanging the elements accordingly. The heuristic terminates when no feasible improving element remains in the "swap" neighborhood. Fisher et al. [38] show that a locally optimal solution obtained using the interchange heuristic is a 2-approximate solution to $f_S|\mathcal{F}_P$ (and more generally, over arbitrary matroids). Any locally optimal solution in a "swap" neighborhood to $f_S|\mathcal{F}_P$ in fact corresponds to a pure-strategy Nash equilibrium in a basic-utility game. The result of Fisher et al. [38] would imply that any pure-strategy Nash equilibrium in a basic utility game is at least half of the social optimal value. The result of Vetta implied in Corollary 2.6 is more general, in that it holds for valid utility games and for mixed strategy Nash equilibria as well, although the structure of both proofs are similar in spirit.

Mirrokni and Vetta [75] also consider the notion of a state graph $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ corresponding to a utility game, where each vertex in $\mathcal{V}$ represents a strategy state, $S = (s_1, \ldots, s_n)$. There is a directed edge in $\mathcal{E}$ from state S to S' with label j if the only difference between S and S' is the strategy of player j; and player j plays its best response in strategy state S to go to S'. A one-round best-response path is a path P that starts from an arbitrary state and the edges of P are labeled in order $i_1, i_2, \ldots, i_n$, where $i_1, i_2, \ldots, i_n$ is an arbitrary ordering of

the n players. One may easily define a similar state graph and related notions for a partition matroid game.

We claim that starting with an initial state $(\emptyset_1, \emptyset_2, \ldots, \emptyset_n)$ in the state graph of a basic utility game and following a one-round best response path would correspond to the execution of the locally greedy algorithm of the underlying partition matroid. To see this, without loss of generality, one may assume that the best response path in consideration is labeled $1, 2, \ldots, n$. Furthermore, let the vertices in this path correspond to $S_0 = (\emptyset_1, \emptyset_2, \ldots, \emptyset_n)$, $S_1 = (s_1, \emptyset_2, \ldots, \emptyset_n)$, $\ldots$, $S_n = (s_1, s_2, \ldots, s_n)$ in order. Now, in any iteration j of the locally greedy algorithm, the role of the incremental oracle is to pick an element $e \in E_j$ of maximum possible incremental value, $\rho_e(F_{j-1}) = f(F_{j-1} \cup e) - f(F_{j-1})$, to the current solution at hand, F_{j-1}. In the one-round best response path being considered, the social objective at vertex S_{j-1} is given by $g(\cup_{i=1}^{j-1} s_i)$. By induction, suppose that $F_{j-1} = \{e_i | i = 1, \ldots, j-1\}$ where each e_i corresponds to the strategy s_i of player i in state $S_{j-1} = (s_1, \ldots, s_{j-1}, \emptyset_j, \ldots, \emptyset_n)$. Clearly by definition,

$$f(F_{j-1}) = g(\cup_{i=1}^{j-1} s_i) = \gamma(S_{j-1}).$$

Moreover, in transitioning from S_{j-1} to S_j, player j selects s so as the maximize

$$\alpha_j(T) = g(\cup_{i=1}^{j-1} s_j \cup s) - g(\cup_{i=1}^{j-1} s_j),$$

where $T = (s_1, \ldots, s_{j-1}, s, \ldots, \emptyset)$. However, observe that

$$\alpha_j(T) = f(F_{j-1} \cup e) - f(F_{j-1}) = \rho_e(F_{j-1}),$$

where e would be the element in E_j corresponding to the strategy s. Hence, it must be that the element selected by the locally greedy algorithm is indeed the e_j that corresponds to s_j. The claim follows by induction.

Mirrokni and Vetta [75] show that a one-round best response path starting from the initial state, $(\emptyset_1, \emptyset_2, \ldots, \emptyset_n)$, provides a 2-approximation to the state S that maximizes $\gamma(S)$. By our interpretation of this path as the execution of the locally greedy algorithm, the same

result would follow from Theorem 2.5.

In thinking about partition matroid games, an interesting question that arises if one might be able to imbue players with private valuation functions so that good approximate solutions may be found for $f_S | \mathcal{F}_P$ using iterated best responses, or even improved responses, as discussed earlier. Alternately posed, recall that in an underlying state graph, directed edges between vertices are determined by the private utility functions of players, and the corresponding best responses based on these utility functions. Therefore, by appropriately selecting the players' private utility functions, may one reorient edges in the graph so that starting from any state, one is guaranteed to reach a state with a good social optimal value in a polynomial number of steps? Goemans et al. [45] show that indeed this would be possible by setting up the players utilities as a basic-utility game. Specifically, they show that:

Theorem 2.7. (Goemans et al. [45]) *In basic-utility games, for any constant $\epsilon > 0$, there exists a constant c such that the expected social value of a state after $cn \log \frac{1}{\epsilon}$ random best responses is at least $\frac{1}{2} - \epsilon$ of the optimum. Moreover, for any constant $\epsilon' > 0$, there exists constants $\epsilon, c' > 0$ such that after $c'n \log n \log \frac{1}{\epsilon}$ random best responses, the social value is at least $\frac{1}{2} - \epsilon'$ of the optimum with high probability.*

2.5.5 Set k-Cover Problems in Wireless Sensor Networks

Motivated by applications in wireless sensor networks, Abrams et al. [2] consider the following variant of the SET k-COVER problem:

SET k-COVER [2]

Instance: A set of elements, U, a collection S of subsets of U, and an integer $k \geq 2$.

Task: Find a partition of the collection of subsets S into k parts, $C_1, \ldots, C_k$ such that $\sum_{i=1}^{k} | \cup_{S_j \in C_i} S_j |$ is maximized.

The intuition behind the formulation of this problem is as follows: the underlying elements of the set U are meant to represent distinct regions being monitored by a sensor network, and each subset $S_i \in S$ represents the regions monitored by a particular wireless sensor i.

75

The objective of the planner is to partition these sensors into k parts so as to maximize the number of times the regions are covered by these parts. Each part of the partition corresponds to a group of sensors that are activated for a particular period of time, and different parts of the partition are activated at different times, so as to conserve the battery power of the wireless sensors.

Consider a ground set $E = \cup_{j=1}^{|\mathcal{S}|} E_j$, where any element $e_{ij} \in E_j$ corresponds to assigning set S_j to partition $\mathcal{C}_i$. That a set in $\mathcal{S}$ may be allocated to at most one partition defines a partition matroid on E. Moreover, for any subset $F \subseteq E$, create a partition with $\mathcal{C}_i^F = \{S_j | e_{ij} \in F\}$. Now, the objective function of the SET k-COVER problem would correspond to:

$$f(F) = \sum_{i=1}^{k} | \cup_{S_j \in \mathcal{C}_i^F} S_j|.$$

It is not hard to see that f is nondecreasing and submodular, using a similar argument as seen for MAXIMUM COVERAGE. Therefore, this problem is an instance of $f_S | \mathcal{F}_P$. Abrams et al. [2] propose a number of algorithms, including a distributed greedy algorithm, and show that it is a 2-approximation algorithm for the problem. This distributed greedy algorithm is in fact analogous to the locally greedy algorithm, and the performance of the distributed greedy algorithm of Abrams et al. [2] follows from Theorem 2.5.

2.5.6 Applications of Generalized Result for Partition Matroids

Based on the result of Theorem 2.5, we now put into perspective other results in the literature where in the absence of an optimal incremental oracle, the locally greedy algorithm uses an α-approximate incremental oracle. As mentioned earlier, the SEPARABLE ASSIGNMENT problem is an instance of $f_S | \mathcal{F}_P$. Fleischer et al. [39] devise a polynomial-time local search $(\alpha+1+\epsilon)$-approximation algorithm for SEPARABLE ASSIGNMENT, given an α-approximation algorithm for the single-bin subproblem. It may be seen that any such α-approximation algorithm for the single-bin subproblem corresponds exactly to an α-approximate incremental oracle for the locally greedy algorithm. Theorem 2.5 therefore implies that:

Corollary 2.8. *There is a polynomial-time locally greedy $(\alpha + 1)$-approximation algorithm for* SEPARABLE ASSIGNMENT, *given an α-approximation algorithm for the single-bin subproblem.*

Fleischer et al. [39] do propose a linear programming-based $\alpha e/(e-1)$-approximation algorithm for SEPARABLE ASSIGNMENT, given an α-approximation algorithm for the single-bin subproblem. However, observe that if $\alpha \geq (e-1)$, then $(\alpha + 1) \leq \alpha e/(e-1)$. Hence, if we only have "weak" approximation algorithms for the single-bin subproblem (such as the RECTANGLE PACKING problem), the locally greedy algorithm outperforms the LP-based algorithm for SEPARABLE ASSIGNMENT.

Chekuri and Khanna [15] prove that for the MULTIPLE KNAPSACK problem, the performance ratio of a greedy algorithm solving the KNAPSACK problem successively is $(2 + \epsilon)$, and note that the same result holds even when the weights of items vary across bins. Also, Dawande et al. [25] propose a similar greedy algorithm-based $(2 + \epsilon)$-result for a MULTIPLE KNAPSACK problem with "assignment restrictions," wherein items are restricted to be assigned only to certain specified sets of bins. We note that both these results follow from the corollary above, since these problems are special cases of SEPARABLE ASSIGNMENT.

Chekuri and Kumar [16] study a variant of the MAXIMUM COVERAGE problem, that they call MAXIMUM COVERAGE WITH GROUP BUDGET CONSTRAINTS. They also consider the performance of a greedy algorithm that uses an α-approximate incremental oracle and show that the performance of their greedy algorithm for the cardinality version of MAXIMUM COVERAGE WITH GROUP BUDGET CONSTRAINTS is $(\alpha + 1)$. By observing that the cardinality version of MAXIMUM COVERAGE WITH GROUP BUDGET CONSTRAINTS is a special case of $f_S|\mathcal{F}_P$, and that their greedy algorithm is analogous to the locally greedy algorithm, Theorem 2.5 implies the same result.

In Section 1.9, we discussed the ADWORDS ASSIGNMENT problem introduced by Fleischer et al. [39], and showed in Lemma 1.11 that it was a special case of $f_S|\mathcal{F}_P$. From the transformation presented, it is not hard to see that the role of the incremental oracle in a locally greedy algorithm for this problem would be played by an algorithm for the RECTANGLE

PACKING problem discussed in Section 1.2.3. For RECTANGLE PACKING, the best-known result is a $(2+\epsilon)$-approximation scheme due to Jansen and Zhang [58]. Consequently, Theorem 2.5 would imply that:

Theorem 2.9. *The locally greedy algorithm, with a $(2+\epsilon)$-approximation scheme for* RECTANGLE PACKING *as an approximate incremental oracle, is a $(3 + \epsilon)$-approximation scheme for the* ADWORDS ASSIGNMENT *problem.*

The above result improves on the previous best $(2 + \epsilon)e/(e - 1) \approx (3.16 + \epsilon)$-approximation result of Fleischer et al. [39] for the ADWORDS ASSIGNMENT problem.

It is instructive to understand the difference between the locally greedy algorithm and the standard greedy algorithm. A standard greedy algorithm in any iteration tries to pick the "best" incremental element in E over all element types, and does not constrain itself to pick only from a certain subset E_i. One might therefore expect that for partition matroids, the standard greedy performs better than the locally greedy algorithm. However, as it turns out, even the standard greedy algorithm achieves the same approximation factor as the locally greedy algorithm, and this factor is tight.

Observation 2.10. *For the problem $f_S|\mathcal{F}_P$, the worst-case performance of a standard greedy algorithm as well as a locally greedy algorithm using an α-approximate incremental oracle is no better than $(\alpha + 1)$.*

Proof. Consider a partition matroid, with $E = E_1 \cup E_2$, with $E_1 = \{a, b\}, E_2 = \{c\}$, where at most one element may be picked from E_1 and E_2 respectively, and a submodular function, f defined as $f(\emptyset) = 0, f(\{a\}) = \alpha, f(\{b\}) = f(\{c\}) = 1, f(\{a,c\}) = f(\{a,b\}) = \alpha + 1, f(\{b,c\}) = 1, f(\{a,b,c\}) = \alpha + 1$. It may be easily verified that f is indeed normalized, nondecreasing and submodular. Moreover, the optimal solution in this instance yields a value of $f(\{a,c\}) = \alpha + 1$. However, the standard greedy algorithm and the locally greedy algorithm may yield the solution $f(\{b,c\}) = 1$, by picking b in the first iteration, using an α-approximate incremental oracle. Thus the approximation guarantees of both these algorithms is $(\alpha + 1)$. ∎

We end this section by pointing out that indeed, for a lot of the problems discussed, there are algorithms with better performance guarantees than that of the locally greedy algorithm. For example, Fleischer et al. [39] give a $(e/(e-1))$-approximation algorithm for the GENERALIZED ASSIGNMENT problem, that has recently been improved to $(e/(e-1)-\epsilon)$ by Feige and Vondrák [36]. For the MULTIPLE KNAPSACK problem, there exists a PTAS constructed by Chekuri and Khanna [15]. Abrams et al. [2] also provide an $(e/(e-1))$-approximation algorithm for their variant of SET k-COVER. For the winner determination problem in combinatorial auctions with submodular valuations, Feige and Vondrák [36] develop a $(e/(e-1)-\epsilon)$-approximation algorithm using a *demand oracle*, and Dobzinski and Schapira [31] develop an $(2-1/n)$-approximation algorithm using a *value oracle*. A demand oracle, given a set of prices p_i, one for each element e_i, outputs a set S that maximizes $f(S) - \sum_{e_i \in S} p_i$. This assumption of the existence of a demand oracle is a stronger assumption than that of a value oracle, as it may shown that a demand oracle simulates a value oracle in polynomial time (Dobzinski and Schapira [31]). Recall that the existence of a value oracle is sufficient to show the factor-2 performance of a locally greedy algorithm for this problem. Nevertheless, these results do hint that there might be better approximation algorithms for the general class of problems, $f_S | \mathcal{F}_P$, itself.

Calinescu et al. [14] have recently developed an $(e/(e-1)$-approximation algorithm for the problem of maximizing the sum of weighted rank functions of matroids over an arbitrary matroid constraint. The sum of weighted rank functions are a rich subclass of monotone submodular functions, and include most of the objective functions of problem instances discussed in this chapter, but there do exist instances of monotone submodular functions that do not belong in the above class, notably including the objective function illustrated for the ADWORDS ASSIGNMENT problem in Lemma 1.11. Calinescu et al. [14] also conjecture that an $e/(e-1)$-approximation algorithm exists for maximizing monotone submodular functions over arbitrary matroids, $f_S | \mathcal{F}_M$. In the following section, we show if the size of the ground set, E, is polynomial in the input size, then there is a randomized $(2-1/n)$-approximation algorithm for $f_S | \mathcal{F}_P$, where $n = \max_i |E_i|$.

2.6. An Improved Randomized Algorithm for Partition Matroids

In Section 2.5.3, we showed that the WINNER DETERMINATION problem in combinatorial auctions with submodular bidders is in fact a special case of $f_S|\mathcal{F}_P$. We also observed that the objective of the WINNER DETERMINATION problem, while submodular, has a special *separable* structure. More specifically, in the WINNER DETERMINATION problem posed as $f_S|\mathcal{F}_P$, the items assigned to bidder j do not affect the resulting valuation of player i, whereas in arbitrary submodular functions defined over ground sets, this might well be the case. Dobzinski and Schapira [31] propose a randomized algorithm that guarantees a $(2 - \frac{1}{n})$-approximate solution in expectation for the WINNER DETERMINATION problem, where n is the number of bidders in the auction. In this section, we adapt the algorithm of Dobzinski and Schapira [31] for the more general problem of $f_S|\mathcal{F}_P$ and show that it is a $(2 - \frac{1}{n})$-approximation algorithm for this problem as well, where $n = \max_i |E_i|$.

Recall that we had earlier defined a partition matroid, $(E, \mathcal{F})$, over a ground set, $E = \cup_{i=1}^{m} E_i$ as:

$$\mathcal{F} = \{F : F = \cup_{i=1}^{m} F_i \text{ where } F_i \subseteq E_i, |F_i| \leq l_i \text{ for } i = 1, \ldots, m\}.$$

If the objective is to maximize a nondecreasing submodular set function f over this partition matroid, then we claim that without loss of generality, it is sufficient to consider the case when at most one element may be picked of each type, E_i. Instead of picking l_i elements from E_i, we could make l_i copies of E_i and pick at most one element from each. Moreover, to ensure that duplicate copies of the same item are not selected in a solution, we extend the original nondecreasing submodular function, f, so that the incremental value of an element e to a set S is 0 if a duplicate of element e is already present in S. If no duplicate of element e is present in S, then the incremental value of e to S only depends on the distinct underlying elements in S. By this definition of the new objective function, it is clearly nondecreasing, since adding new elements can only increase its value. Moreover, the incremental value of

any element in the new partition matroid also decreases as the underlying set grows, and therefore the new objective function still remains submodular.

Observe that the duplication process described above is "efficient," in the sense that any optimal solution must contain $O(l_i)$ elements of each type (without loss of generality, $l_i \leq |E_i|$), and in our transformation, we multiply the size of the input by at most $\max_i l_i$. Thus, it is sufficient to consider the following definition of a partition matroid, $(E, \mathcal{F})$, for our purposes:

$$E = \cup_{i=1}^m E_i \quad \text{and} \quad \mathcal{F} = \{F \subseteq E : |F \cap E_i| \leq 1 \quad \text{for } i = 1, \ldots, m\}$$

Recall that $n = \max_i |E_i|$. Without loss of generality, we may also assume that each set E_i has n elements, by adding dummy elements of 0 incremental value if E_i has fewer than n elements. In addition, we extend the submodular function so that the incremental value of any of the original items does not depend on any of the dummy elements in the underlying set. Using similar ideas as above, it may again be verified that the modified objective function is still nondecreasing submodular. Note that the process of adding dummy elements to E_i is guaranteed to be polynomial if the ground set size, $|E|$, is polynomial in the size of the input. By arbitrarily indexing elements, we further assume that $E_i = \{e_{ij} | j = 1, \ldots, n\}$ for all $i = 1, \ldots, m$.

We are now ready to present the randomized algorithm:

RANDOMIZED ALGORITHM FOR $f_S | \mathcal{F}_P$

Step 1: Set $i := 1$; let $S_0 := \emptyset$.

Step 2: For all elements of type i:

Let $t_j = (f(S_{i-1} \cup \{e_{ij}\}) - f(S_{i-1}))^{n-1}$ for each $j = 1, \ldots, n$.

Select e_{ij} with probability $q_j = \frac{t_j}{\sum_{k=1}^n t_k}$.

Step 3: Set $i := i + 1$, $S_i := S_{i-1} \cup \{e_{ij}\}$. If $i \leq m$, then goto **Step 2**.

Intuitively, the randomized algorithm in each iteration, picks an element of a particular type with a probability that depends on its incremental value to the solution relative to the

incremental value of other elements of the same type. However, the algorithm scales the relative impact of each element appropriately by raising the incremental value of an element to the power of $n - 1$. We shall discuss the potential impact of other scaling measures later in the section.

Observe that the algorithm presented above requires access to a value oracle that given a set S, returns the value $f(S)$. Recall that the greedy algorithm and the locally greedy algorithm only required access to an incremental oracle. A value oracle can simulate an incremental oracle by enumerating the incremental value of all elements and then selecting the element with the best incremental value. Hence, the assumption of a value oracle is a stronger than that of an incremental oracle. However, for most applications, including those discussed in this chapter, this is a reasonable assumption.

Observation 2.11. *The* RANDOMIZED ALGORITHM FOR $f_S|\mathcal{F}_P$ *requires the existence of a value oracle, a stronger assumption than an incremental oracle.*

The reader will also note that the running time of the RANDOMIZED ALGORITHM FOR $f_S|\mathcal{F}_P$ depends polynomially on n, the number of element types, m, and the running time of the value oracle. In fact, the algorithm runs in time $O(n^2 m + nm\mathrm{P})$, where P is the running time of the value oracle, since the value oracle is invoked once for each element, and it takes $O(n)$ time to compute the probability of picking each element, with there being at most nm elements in E.

Our main performance result regarding the algorithm is that:

Theorem 2.12. *The* RANDOMIZED ALGORITHM FOR $f_S|\mathcal{F}_P$ *is a* $(2 - \frac{1}{n})$-*approximation algorithm.*

Proof. We prove the result by induction on the number of element types, m.

Let $m = 1$, i.e., there are elements of only one type, $E = E_1$, and we may pick at most one of them. Let I denote an instance of this problem. Furthermore, let $\rho_j = f(\{e_{1j}\})$ denote the value of each element. Clearly, the optimal solution of this problem is $OPT(I) = \rho_k = \max_{j=1,\dots,n} \rho_j$. Moreover, the expected value of the solution returned by the randomized

algorithm, $E[ALG(I)]$, is given by:

$$E[ALG(I)] = \sum_{j=1}^{n} q_j \rho_j = \sum_{j=1}^{n} \frac{\rho_j^n}{\sum_{i=1}^{n} \rho_i^{n-1}} = \frac{\sum_j \rho_j^n}{\sum_j \rho_j^{n-1}}.$$

Recall that q_j denotes the probability that the RANDOMIZED ALGORITHM FOR $f_S|\mathcal{F}_P$ picks element e_{1j}. We would like to verify that $OPT(I) \leq (2-\frac{1}{n})E[ALG(I)]$. From the expressions for $OPT(I)$ and $E[ALG(I)]$ above, we have that:

$$\frac{OPT(I)}{E[ALG(I)]} = \frac{\rho_k \sum_j \rho_j^{n-1}}{\rho_k^n + \sum_{j \neq k} \rho_j^n} = \frac{\rho_k^n + \rho_k \sum_{j \neq k} \rho_j^{n-1}}{\rho_k^n + \sum_{j \neq k} \rho_j^n} \leq 1 + \frac{\rho_k \sum_{j \neq k} \rho_j^{n-1}}{\rho_k^n + \sum_{j \neq k} \rho_j^n}. \tag{2.9}$$

Let $\bar{\rho}$ denote $\frac{\sum_{j \neq k} \rho_j^{n-1}}{n-1}$. Then, it follows from the power means inequality (recall inequality (2.1) from Section 2.1.1) that:

$$\left(\frac{\sum_{j \neq k} \rho_j^n}{n-1} \right)^{1/n} \geq \bar{\rho}^{\frac{1}{n-1}}. \quad \text{Hence,} \quad \sum_{j \neq k} \rho_j^n \geq (n-1)\bar{\rho}^{\frac{n}{n-1}}. \tag{2.10}$$

Substituting $\bar{\rho}$ and inequality (2.10) into inequality (2.9), we have that:

$$\frac{OPT(I)}{E[ALG(I)]} \leq 1 + \frac{\rho_k(n-1)\bar{\rho}}{\rho_k^n + (n-1)\bar{\rho}^{\frac{n}{n-1}}}. \tag{2.11}$$

Using elementary calculus, it may be verified that the expression in the right-hand side of inequality (2.11) is concave in $\bar{\rho}$ and maximized when $\bar{\rho} = \rho_k^{n-1}$. This would yield that:

$$\frac{OPT(I)}{E[ALG(I)]} \leq 1 + \frac{\rho_k(n-1)\rho_k^{n-1}}{\rho_k^n + (n-1)\rho_k^n} = 1 + \frac{n-1}{n} = 2 - \frac{1}{n},$$

thus proving the claim for $m = 1$.

Assuming that the claim holds for $m - 1$, we seek to show that the claim holds for m. Suppose again that I denotes an instance of the problem with m types of items. Additionally, let I^j denote the instance of the problem where element e_{1j} is selected among all items of type 1. Thereby, in I^j, one would need to select only among the elements of $m - 1$ types,

$E_2, \ldots, E_m$. In order for I^j to be an instance of $f_S|\mathcal{F}_P$, the objective function, f, in I^j must be normalized. We may do this by redefining the objective of I^j as:

$$f^j(S) = f(S \cup \{e_{1j}\}) - f(\{e_{1j}\}) \text{ for all } S \subseteq \cup_{i=2}^m E_i.$$

By again denoting $\rho_j = f(\{e_{1j}\})$, we would then have that $f(S \cup \{e_{1j}\}) = f^j(S) + \rho_j$.

Let the random variable $ALG(I)$ denote the value of a solution generated by the randomized algorithm for instance I. Similarly, let $ALG(I^j)$ denote the value of the solution generated by the randomized algorithm for the instance I^j. Then, it follows from the description of the algorithm that:

$$E[ALG(I)] = \sum_{j=1}^n q_j \left(E[ALG(I^j)] + \rho_j \right). \tag{2.12}$$

Moreover, by the induction assumption, we also have that $OPT(I^j) \leq (2 - \frac{1}{n})E[ALG(I^j)]$. If we had the inequality that:

$$OPT(I) \leq \sum_{j=1}^n q_j OPT(I^j) + (2 - \frac{1}{n}) \sum_{j=1}^n \rho_j q_j, \tag{2.13}$$

then it would follow using the induction assumption and inequality (2.12) that:

$$\begin{aligned} OPT(I) &\leq (2 - \frac{1}{n}) \sum_{j=1}^n q_j E[ALG(I^j)] + (2 - \frac{1}{n}) \sum_{j=1}^n \rho_j q_j \\ &= (2 - \frac{1}{n}) \sum_{j=1}^n q_j \{E[ALG(I^j)] + \rho_j\} = (2 - \frac{1}{n})E[ALG(I)], \end{aligned}$$

thus proving the induction claim. Hence, we are left to show that inequality (2.13) holds.

Lemma 2.13. $\sum_{j=1}^n q_j \left(OPT(I) - OPT(I^j) \right) \leq (2 - \frac{1}{n}) \sum_{j=1}^n \rho_j q_j$.

Let $O \subseteq E$ denote an optimal solution to instance I. Also, among elements of type 1, let $e_{1k} \in O$ and denote $O' = O \setminus \{e_{1k}\}$. Then, O' is a feasible solution to each instance, I^j, since it does not contain any element of type 1. To prove the claim, we bound the value of

84

$OPT(I) - OPT(I^j)$ for each $j = 1, \ldots, n$.

Observe that O' must be an optimal solution to instance I^k. If it were not, and S were instead the optimal solution to I^k, then observe that by adding element e_{1k} to set S, we would have a feasible solution to I with objective value:

$$f(S \cup \{e_{1k}\}) = f^k(S) + \rho_k > f^k(O') + \rho_k = f(O' \cup \{e_{1k}\}) = f(O) = OPT(I),$$

where the second inequality follows from the assumption that O' is not an optimal solution to I^k. This is a contradiction, since O is the optimal solution to instance I. Now since O' is an optimal solution to I^k, it follows that:

$$OPT(I) - OPT(I^k) = f(O) - f^k(O') = \rho_k. \tag{2.14}$$

For any $j \neq k$, it follows that since O' is a feasible solution to instance I^j:

$$
\begin{align}
OPT(I) - OPT(I^j) \quad &\leq \quad f(O) - f^j(O') \tag{2.15a} \\
&= \quad f(O) - (f(O' \cup \{e_{1j}\}) - \rho_j) \tag{2.15b} \\
&\leq \quad f(O') + \rho_k - f(O' \cup \{e_{1j}\}) + \rho_k \tag{2.15c} \\
&\leq \quad \rho_k + \rho_j \ . \tag{2.15d}
\end{align}
$$

Inequality (2.15a) follows from the fact that $f^j(O') \leq OPT(I^j)$; inequality (2.15b) follows from the definition of $f^j(.)$; inequality (2.15c) from the submodularity of f, and inequality (2.15d) follows since f is nondecreasing.

From equations (2.14) and (2.15d), and since $q_j = \frac{\rho_j^{n-1}}{\sum_i \rho_i^{n-1}}$ we now have that:

$$
\begin{align}
\sum_{j=1}^{n} q_j \left(OPT(I) - OPT(I^j) \right) \quad &= \quad \frac{\sum_{j=1}^{n} \rho_j^{n-1} \left(OPT(I) - OPT(I^j) \right)}{\sum_{j=1}^{n} \rho_j^{n-1}} \\
&\leq \quad \frac{\left[\sum_{j \neq k} \rho_j^{n-1} (\rho_j + \rho_k) \right] + \rho_k^n}{\sum_{j=1}^{n} \rho_j^{n-1}} \ .
\end{align}
$$

85

Moreover, dividing both sides by $\sum_{j=1}^{n} \rho_j q_j$, we have that:

$$\frac{\sum_{j=1}^{n} q_j \left(OPT(I) - OPT(I^j) \right)}{\sum_{j=1}^{n} \rho_j q_j} \leq \frac{\frac{\left[\sum_{j \neq k} \rho_j^{n-1} (\rho_j + \rho_k) \right] + \rho_k^n}{\sum_{j=1}^{n} \rho_j^{n-1}}}{\sum_{j=1}^{n} \rho_j q_j}$$

$$= \frac{\frac{\left[\sum_{j \neq k} \rho_j^{n-1} (\rho_j + \rho_k) \right] + \rho_k^n}{\sum_{j=1}^{n} \rho_j^{n-1}}}{\frac{\sum_{j=1}^{n} \rho_j^n}{\sum_{j=1}^{n} \rho_j^{n-1}}} = \frac{\left[\sum_{j \neq k} \rho_j^{n-1} (\rho_j + \rho_k) \right] + \rho_k^n}{\sum_{j=1}^{n} \rho_j^n}$$

$$= 1 + \frac{\rho_k \sum_{j \neq k} \rho_j^{n-1}}{\rho_k^n + \sum_{j \neq k} \rho_j^n} \,. \qquad (2.16)$$

Observe that the structure of the right-hand side of equation (2.16) exactly matches that of equation (2.9). In a similar manner as derived for equation (2.9), it therefore follows that:

$$\frac{\sum_{j=1}^{n} q_j \left(OPT(I) - OPT(I^j) \right)}{\sum_{j=1}^{n} \rho_j q_j} \leq \left(2 - \frac{1}{n} \right),$$

thereby completing the proof. ∎

We return now to a question that we had posed earlier with respect to **Step 2** of the RANDOMIZED ALGORITHM. What if one considers the modified incremental value of each element, $t_j = (f(S_{i-1} \cup \{e_{ij}\}) - f(S_{i-1}))^p$, for some arbitrary integer p ? In fact, the randomized algorithm would behave asymptotically like the a *randomized version* of the locally greedy algorithm as $p \to \infty$, since the probability of picking the element with the largest incremental value (if unique) in each iteration would approach 1, or else the probability would be equally distributed amongst all elements with the highest incremental value. It is not hard to show that if $p > 0$, the analysis of Theorem 2.12 would reveal that this algorithm is a $\left(1 + \left(\frac{p}{n-1} \right)^{\frac{p}{p+1}} \frac{n-1}{p+1} \right)$-approximation algorithm. Moreover, using elementary calculus, it may be shown that this factor is indeed minimized when $p = n - 1$. In the case that $p \leq 0$, the analysis does not hold true any longer.

When $p \leq 0$, the algorithm does not make intuitive sense, since items with a relatively high incremental value are picked with a lower relative probability. We demonstrate this via an example.

86

Example 2.14. *Consider a uniform matroid, a special case of the partition matroid, with $E = \{e_1, e_2, \ldots, e_n\}$ and a modular function, f defined as $f(e_1) = 10, f(e_2) = f(e_3) = \ldots = f(e_n) = 1$. The optimal solution to this problem is $OPT(I) = 10$. If $p \leq 0$, then each element $e_i, i \geq 2$ is picked with probability at least $\frac{1}{n}$, and element e_1 is picked with probability at most $\frac{1}{n}$ by the* RANDOMIZED ALGORITHM FOR $f_S|\mathcal{F}_P$. *This would yield that $ALG(I) \leq \frac{n+9}{n}$, making the performance of the algorithm at least 2 is $n \geq 3$.*

Thus in the presence of a value oracle, the ability to randomize over elements rather than acting in a locally optimal manner helps achieve a better performance ratio for $f_S|\mathcal{F}_P$. Combining aspects of the randomized and greedy approaches might pave the path towards closing the approximation gap between $(2 - \frac{1}{n})$ and $(\frac{e}{e-1})$, the best known lower bound on approximability for the problem, unless P=NP.

As mentioned earlier, as $p \to \infty$, the RANDOMIZED ALGORITHM FOR $f_S|\mathcal{F}_P$ behaves as a natural randomized extension of the greedy algorithm. Given a value oracle, this extension to the greedy algorithm would be that if in any iteration, more than one element have the same best incremental value, then the algorithm picks any one of these "best elements" with equal probability. However, as we show in the example below, this randomized greedy algorithm would also have a worst-case guarantee arbitrarily close to 2.

Example 2.15. *Consider the partition matroid specified by $E = E_1 \cup E_2$, where $E_1 = \{a, b\}$ and $E_2 = \{c\}$ and at most one element may be picked of each type. Furthermore, consider the submodular function defined by: $f(\emptyset) = 0, f(\{a\}) = 1, f(\{b\}) = 1 + \epsilon, f(\{c\}) = 1, f(\{a, c\}) = 2, f(\{a, b\}) = 2, f(\{b, c\}) = 1 + \epsilon, f(\{a, b, c\}) = 2$. It may be easily verified that f is indeed normalized, nondecreasing and submodular. Moreover, the optimal solution in this instance yields a value of $OPT(I) = f(\{a, c\}) = 2$. However, the randomized versions of the greedy algorithm and the locally greedy algorithm behave exactly as the standard versions, and may yield the solution $f(\{b, c\}) = 1 + \epsilon$. Thus the approximation guarantees of the randomized versions of these algorithms approach 2 as $\epsilon \to 0$.*

2.6.1 Application to Max Sat

Recall from the transformation of MAX SAT to $f_S|\mathcal{F}_P$ presented in Section 2.5.2, that for each variable x_i, $|E_i| = 2$, implying therefore that $n = 2$. Additionally, since the number of variables is polynomial in the input size of MAX SAT, it also follows that $|E|$ is polynomial in the input size of MAX SAT. This would make MAX SAT an ideal candidate problem for which the RANDOMIZED ALGORITHM FOR $f_S|\mathcal{F}_P$ would be a polynomial time approximation algorithm. Therefore, from Theorem 2.12, it follows that:

Corollary 2.16. *The* RANDOMIZED ALGORITHM FOR $f_S|\mathcal{F}_P$ *is a 3/2-approximation algorithm for* MAX SAT.

Interestingly, this performance matches the tight 3/2-factor performance of Johnson's algorithm for MAX SAT (Johnson [60]), as shown by Chen et al. [17]. Whether the performance of RANDOMIZED ALGORITHM FOR $f_S|\mathcal{F}_P$ is tight for MAX SAT, or even whether the analysis presented in Theorem 2.12 is tight, remains an open question.

2.7. Generalized Results over Matroids and Independence Systems

Observe that for uniform matroids and partition matroids, the admissibility[1] of an element into an independent set S depends only on the number of elements of each type present in S. Since the admissibility of an element into the greedy solution can be determined trivially for partition matroids and does not involve the need of an admissibility oracle, as would be the case for general matroids and independence systems, the above study was simple, and the running time of the greedy algorithms was independent of the size of the ground set, E. If a polynomial-time admissibility oracle does exist for a particular class of matroids or independence systems, then it is possible to study the performance of a greedy algorithm with an α-approximate incremental oracle for such a class of matroids.

---

[1]An element e is said to be "admissible" into an independent set S if $S \cup \{e\}$ remains independent.

Suppose that an independence system $(E, \mathcal{F})$ is the intersection of M different matroids. In this section, we shall generalize the result of Fisher et al. [38], who proved that if an independence system $(E, \mathcal{F})$ is an intersection of a finite number, M, of matroids, then the standard greedy algorithm is a $(M + 1)$-approximation algorithm. More formally, for the problem of maximizing a nondecreasing submodular function over $(E, \mathcal{F})$, we shall show that a greedy algorithm, with an α-approximate incremental oracle as well as an admissibility oracle for $(E, \mathcal{F})$ at its disposal, is in fact an $(\alpha M + 1)$-approximation algorithm.

We begin with a description of a generic greedy algorithm for this problem.

GREEDY ALGORITHM FOR $f_S|\mathcal{F}$

Initialization: Set $i = 1$; let $S_0 = \emptyset$, $E_0 = E$.

Step 1: If $E_{i-1} = \emptyset$, STOP.

Step 2: Select an element $e_i \in E_{i-1}$ for which $\alpha \cdot \rho_{e_i}(S_{i-1}) \geq \max_{e \in E_{i-1}} \rho_e(S_{i-1})$ using an α-approximate incremental oracle.

Step 3: Using the admissibility oracle, check if $S_{i-1} \cup \{e_i\} \in \mathcal{F}$.

Step 4a: If "no," set $E_{i-1} := E_{i-1} \setminus \{e_i\}$ and return to Step 1.

Step 4b: Set $S_i := S_{i-1} \cup \{e_i\}$, $\rho_{i-1} := \rho_{e_i}(S_{i-1})$ and $E_i := E_{i-1} \setminus \{e_i\}$.

Step 5: Set $i := i + 1$ and return to Step 1.

Similar to the standard greedy algorithm described earlier for uniform matroids, the GREEDY ALGORITHM FOR $f_S|\mathcal{F}$ uses the α-approximate incremental oracle to select a candidate element of "good" incremental value. The algorithm then uses the admissibility oracle to check if the selected element is indeed admissible into the solution at hand, and modifies the ground set and solution set accordingly. Observe from the definition of an independence system that if an element e is not admissible to the candidate solution in an iteration i of the algorithm, then it is never admissible to the candidate solution after iteration i. Therefore, such elements may be removed from the underlying set for future consideration.

Theorem 2.17. *Suppose $(E, \mathcal{F})$ is an independence system that can be expressed as the intersection of a finite number, M, of matroids, and f is a normalized, nondecreasing, submodular function. If z_g is the value of the greedy heuristic solution, utilizing an α-approximate*

incremental oracle and an admissibility oracle, for the following problem:

$$\max\{f(S) : S \in \mathcal{F}\}$$

and z_{opt} is the value of an optimal solution, then $\frac{z_{opt}}{z_g} \leq \alpha M + 1$.

Proof. Let us define U_t to be the set of elements considered in the first $(t + 1)$ iterations of the greedy algorithm before the addition of the $(t + 1)$st element. Let $r_m(S)$ denote the rank of set S in matroid m (where the rank of S is the cardinality of the largest independent subset of S in the matroid), and $sp_m(S)$ be the span of S in matroid m, defined by:

$$sp_m(S) = \{e \in E : r_m(S \cup \{e\}) = r_m(S)\} \ .$$

In order to proceed with the proof, we shall utilize two lemmata shown by Fisher et al. [38].

Lemma 2.18 (Fisher et al. [38]). $U_t \subseteq \bigcup_{m=1}^{M} sp_m(S_t) \quad t = 0, 1, \dots \ .$

Proof. If $j \in U_t$, then either $j \in S_t \subseteq \bigcup_{m=1}^{M} sp_m(S_t)$ for all m, or j is not admissible according to the admissibility oracle, implying that $j \in sp_m(S_t)$ for some matroid m. ∎

Lemma 2.19 (Fisher et al. [38]). *If $\sum_{i=0}^{t-1} x_i \leq t$ for $t = 1, 2, \dots, K$, and $\rho_{i-1} \geq \rho_i$ with $\rho_i, x_i \geq 0$ for $i = 1, \dots, K - 1$ and $\rho_K = 0$, then $\sum_{i=0}^{K-1} \rho_i x_i \leq \sum_{i=0}^{K-1} \rho_i$.*

Proof. Consider the following linear program:

$$V = \max_{x} \left\{ \sum_{i=0}^{K-1} \rho_i x_i \ \Big| \ \sum_{i=0}^{t-1} x_i \leq t, t = 1, \dots, K, \quad x_i \geq 0, i = 0, \dots, K - 1 \right\}.$$

It is easy to verify that its dual is:

$$W = \min_{z} \left\{ \sum_{t=1}^{K} t z_{t-1} \ \Big| \ \sum_{t=i}^{K-1} z_t \geq \rho_i, i = 0, \dots, K - 1, \quad z_t \geq 0, t = 0, \dots, K - 1 \right\}.$$

90

As $\rho_i \geq \rho_{i+1}$, the solution $z_i = \rho_i - \rho_{i+1}, i = 0, \ldots, K - 1$ (where $\rho_K = 0$) is dual feasible with value $\sum_{t=1}^{K} t(\rho_{i-1} - \rho_i) = \sum_{t=0}^{K-1} \rho_i$. By weak LP duality, the result follows. ∎

Suppose that S and T represent the greedy and an optimal solution, respectively, to the above problem. Additionally, let $|S| = K$. Note that since $(E, \mathcal{F})$ is an independence system and not necessarily a matroid, $|T|$ need not be K.

For $t = 1, \ldots, K$, let $s_{t-1} = |T \cap (U_t \setminus U_{t-1})|$, where U_t is the set of elements considered in the first $(t + 1)$ iterations before the addition of a $(t + 1)$st element to S_t. We assume without loss of generality that $U_0 = \emptyset$ and $U_K = E$. Also, let $\rho^*(S_i) = \max_{e \in E_i} \rho_e(S_i)$ for $i = 0, \ldots, K - 1$.

Since f is a nondecreasing submodular set function, Lemma 2.2 yields:

$$z_{opt} = f(T) \leq f(S) + \sum_{e \in T \setminus S} \rho_e(S) . \tag{2.17}$$

Suppose $t \in \{1, 2, \ldots, K\}$ and $\rho_{q(t)} = \min\{\rho_i \mid i = 0, \ldots, t - 1\}$. Now for all elements $e \in T \cap (U_t \setminus U_{t-1})$, we have that:

$$\rho_e(S) \leq \rho_e(S_t) \leq \rho^*(S_t) \leq \rho^*(S_{q(t)+1}) \leq \alpha \rho_{q(t)} . \tag{2.18}$$

While the first inequality follows from the submodularity of f, the second and third follow from the definition of ρ^*. The final inequality follows from the fact that we are using an α-approximate oracle: If e^* is such that $\rho_{e^*}(S_{q(t)+1}) = \rho^*(S_{q(t)+1})$, then by the fact that $e^* \in E_{q(t)+1}$ was not considered by the greedy algorithm in iteration $q(t)$, the inequality follows. Given the above inequality, define $\rho'_{t-1} = \alpha \rho_{q(t)}$. We then have:

$$\rho_e(S) \leq \rho'_{t-1} \quad \text{for all } e \in T \cap (U_t \setminus U_{t-1}), \quad t = 1, 2, \ldots, K. \tag{2.19}$$

Note that by the way that we have defined ρ'_{t-1}, it has the nonincreasing property. In other words, $\rho'_{t-1} \geq \rho'_t$ for all t. This is based on the definition of $q(t)$, which itself has the same nonincreasing property. Additionally, $\rho'_t = \alpha \rho_{q(t+1)} \leq \alpha \rho_t$.

Using this fact in equation (2.17), we now have that:

$$
\begin{aligned}
f(T) &\leq f(S) + \sum_{e \in T \setminus S} \rho_e(S) \\
&\leq f(S) + \sum_{e \in T} \rho_e(S) \\
&= f(S) + \sum_{t=1}^{K} \sum_{e \in T \cap (U_t \setminus U_{t-1})} \rho_e(S) \\
&\leq f(S) + \sum_{t=1}^{K} \rho'_{t-1} s_{t-1}
\end{aligned}
\tag{2.20}
$$

where the last inequality follows from (2.19) and the definition of s_t.

Now, observe that since $s_{t-1} = |T \cap (U_t \setminus U_{t-1})|$, it must be that $\sum_{i=1}^{t} s_{i-1} = |T \cap U_t|$. By Lemma 2.18, we also have that $U_t \subseteq \bigcup_{m=1}^{M} sp_m(S_t)$, which in turn gives us that:

$$
|T \cap U_t| \leq \sum_{m=1}^{M} |T \cap sp_m(S_t)| \ .
$$

But since T is independent in each of the matroids and $r_m(sp_m(S_t)) = t$, it follows that for each m, $|T \cap sp_m(S_t)| \leq t$. This implies that:

$$
\sum_{i=1}^{t} s_{i-1} \leq \sum_{m=1}^{M} |T \cap sp_m(S_t)| \leq M t
\tag{2.21}
$$

where the above inequality is true for each $t = 1, 2, \ldots, K$.

Since $\rho'_t, s_t \geq 0$ for all t, and ρ'_t has the nonincreasing property, by substituting $x_i := s_i$ and $\rho_i := \rho'_i$ in Lemma 2.19, inequality (2.21) now gives us that

$$
\sum_{i=0}^{K-1} \rho'_i s_i \leq M \sum_{i=0}^{K-1} \rho'_i \ .
\tag{2.22}
$$

92

Substituting back inequality (2.22) into inequality (2.20), we now have that:

$$
\begin{aligned}
f(T) &\leq f(S) + M \sum_{i=0}^{K-1} \rho_i' \\
&\leq f(S) + \alpha M \sum_{i=0}^{K-1} \rho_i \qquad\qquad (2.23) \\
&= f(S)(1 + \alpha M)
\end{aligned}
$$

where inequality (2.23) follows from the fact that $\rho_t' \leq \alpha \rho_t$ for all t. ∎

In parallel to this work, Calinescu et al. [14] have also recently noted a somewhat different proof for the performance of the greedy algorithm for $f_S|\mathcal{F}$ in the presence of an α-approximate incremental oracle. In the next section, we present a more substantial discussion of the greedy algorithm and its performance.

2.7.1 Discussion on the Running Time of the Greedy Algorithm

As we have noted in earlier sections for uniform and partition matroids, the greedy algorithm and the locally greedy algorithm have a running time that depends only on the restrictions of the number of elements that must be picked, and is independent of the size of the ground set, E. However, the greedy algorithm presented for $f_S|\mathcal{F}$ in fact has a running time that depends on the size of E. Hence, if E is exponentially large, this would yield a poor running time for the greedy algorithm. It must also be noted that this dependence of the running time on $|E|$ only comes because sometimes, the incremental oracle might pick an element $e \in E$ that need not be admissible. If however, there were a hybrid incremental oracle that always finds a "good" incremental element that is necessarily admissible, then the running time performance of the greedy algorithm will not necessarily depend on $|E|$, but on the size of a largest independent set in E.

In certain instances of $f_S|\mathcal{F}$, such as when the independence system $(E, \mathcal{F})$ is a matroid, it may be possible to bound the size of the largest independent set polynomially in the size of the input of the underlying combinatorial optimization problem. We now present an

93

instance of such a problem, and as it turns out, the greedy algorithm consequently provides the best-known approximation ratio for this problem in polynomial time.

2.7.2 k-Median with Hard Capacities and Packing Constraints

Fleischer et al. [39] present the following variant of the k-Median problem, that they call the k-MEDIAN WITH HARD CAPACITIES AND PACKING CONSTRAINTS problem.

---

k-MEDIAN WITH HARD CAPACITIES AND PACKING CONSTRAINTS [39]

Instance: A set, U, of n items and a set, B, of m bins. Each bin $i \in B$ has an independence system $\mathcal{I}_i$ of subsets of items that fit in bin i. A profit p_{ij} for assigning item j to bin i. An integer $k \leq m$.

Task: Choose a set of K bins, $|K| \leq k$, and a subset of items, $S \subseteq U$, with a feasible assignment of these items to the bins in K, $S_i \in \mathcal{I}_i$ for bin $i \in K$, $S_i \cap S_l = \emptyset$ for $i \neq l$, so as to maximize profit, $\sum_{i \in K} \sum_{j \in S_i} p_{ij}$.

This problem has very similar flavor to the SEPARABLE ASSIGNMENT problem discussed in Section 2.2. In fact, using the transformation of Observation 2.1, it can be seen that the underlying constraint structure is a laminar matroid, defined by $E = \cup_{i=1}^m E_i$, and

$$\mathcal{F} = \{F : F \subseteq E, |F \cap E_i| \leq 1 \text{ for } i = 1, \ldots, m \text{ and } |F| \leq k\} .$$

In the above matroid, it is easy to see that the size of the largest independent set (or basis of the matroid) is k, which is polynomial in the input (since $k \leq m$). Moreover, the objective function for this problem can be rewritten exactly as in Observation 2.1. Consequently, monotone submodularity follows. Thus we establish that K-MEDIAN WITH HARD CAPACITIES AND PACKING CONSTRAINTS is an instance of $f_S|\mathcal{F}_M$.

Fleischer et al. [39] devise a polynomial-time local search ($\alpha + 1 + \epsilon$)-approximation algorithm for K-MEDIAN WITH HARD CAPACITIES AND PACKING CONSTRAINTS, assuming there is an α-approximation algorithm for the single bin subproblem. The authors also

remark that this result is, to the best of their knowledge, the first constant-factor approximation to this problem.

Given an α-approximation algorithm for each of the single bin subproblems corresponding to the m bins, one may easily design an α-approximate *hybrid* incremental oracle over all element types. At the start of any iteration i, suppose that the current solution generated by the greedy algorithm is S. If $|S| \leq k - 1$, then selecting the α-best incremental element among the l element types corresponding to the l bins for which a feasible packing has not been selected as yet, would indeed be a feasible selection. Hence the running time of a greedy algorithm for this problem is polynomial in the input size. Specifically, if P_i is the running time of the α-approximate oracle corresponding to bin i, then the running time of the algorithm is $O(k \sum_{i=1}^{m} P_i)$. Since the problem is an instance of $f_S | \mathcal{F}_M$ with a hybrid incremental oracle available, Theorem 2.17 implies:

Corollary 2.20. *Given an α-approximation algorithm for the single bin subproblem, there is a polynomial-time $(\alpha + 1)$-approximation greedy algorithm for* K-MEDIAN WITH HARD CAPACITIES AND PACKING CONSTRAINTS.

Hence, by generalizing the results of Fisher et al. [38], we are able to improve upon the previous best-known result of Fleischer et al. [39] for this problem. More recently, Calinescu et al. [14] have developed an improved $e/(e - 1)$-approximation algorithm for the GENERALIZED ASSIGNMENT problem subject to a laminar matroid constraint on the bins.

2.8.   Concluding Remarks and Open Questions

In this chapter, we extend some classic results of Fisher et al. [38] and Nemhauser et al. [80] on the performance of the greedy algorithm for maximizing monotone submodular functions over independence systems and other special subclasses. Our work is based on the premise that the greedy algorithm need not always to be able to pick an element of maximum incremental value, and may only be able to select an element of "good" incremental value. We show that this is indeed the case by posing some interesting and important discrete

optimization problems as the problem of maximizing a monotone submodular function over an independence system. Based on our generalized results, we are able to reinterpret as well as present a new view to many recent results. In certain cases, we are even able to establish improved approximation results based on these insights.

We conclude this chapter by highlighting some interesting open questions that remain intimately connected to this work and even served to motivate this study of submodular function maximization:

- Consider the problem $f_S|\mathcal{F}_M$, of maximizing a submodular function over an arbitrary matroid constraint, given a value oracle. Calinescu et al. [14] have recently conjectured that there exists an $e/(e-1)$-approximation problem for this problem. For a rich subclass of submodular functions, Calinescu et al. [14] show that there is indeed such an algorithm, based on pipage rounding and considering an appropriate extension of a submodular function. We believe that the submodular function extension due to Lovász [72] might be of related interest in proving this conjecture using pipage rounding techniques.

- Consider the restricted problem, $f_S|\mathcal{F}_P$, of maximizing a submodular function over a partition matroid. We believe that a insightful first step in proving the conjecture for $f_S|\mathcal{F}_M$ would be to develop an $e/(e-1)$-approximation algorithm to $f_S|\mathcal{F}_P$, given the special simple structure of the partition matroid, that may appeal to developing other combinatorial techniques in proving the conjecture.

- Consider the standard greedy and local greedy algorithms described for $f_S|\mathcal{F}_P$. One would note that in the tight worst-case examples described in Example 2.10, the greedy algorithms' bad performance may be attributed to the greedy algorithm not selecting the "correct" optimal incremental element in the first iteration. While simple randomizing over all optimal incremental elements in any iteration does not necessarily improve the performance of the greedy algorithm (as highlighted in Example 2.15), smarter randomized schemes might lead to improved approximation algorithms for

$f_S|\mathcal{F}_P$. Indeed, we develop an improved randomized $(2-1/n)$-approximation algorithm for $f_S|\mathcal{F}_P$. An interesting question is therefore if one might leverage an intermediate randomized scheme coupled with a greedy strategy to develop improved approximation results for $f_S|\mathcal{F}_P$.

- In Section 2.6, we develop a randomized $(2-1/n)$-approximation algorithm for $f_S|\mathcal{F}_P$, that has a polynomial running time if $|E|$ is polynomial in the input size. However, it remains open whether the analysis presented is indeed tight. Even for the special case of MAX SAT discussed in the section, the result guarantees a 3/2-approximation. However, the worst approximation ratio example that we were able to find was 4/3. It would be interesting to resolve if this analysis is indeed tight, or might be improved.

- In Section 2.5.4, we describe the correspondence between basic-utility games and a decentralized approach to $f_S|\mathcal{F}_P$. We point out an interesting relevant conjecture posed by Mirrokni and Vetta [75] in this context. The authors show that starting from any feasible state in a basic-utility game, a one-round best response path guarantees a state with a 3-approximate social objective value. However, they conjecture that such a path might indeed guarantee a 2-approximate social objective value. Resolving this conjecture is an interesting question as it might suggest a polynomial time 2-approximate local search approach for $f_S|\mathcal{F}_P$.

PART II

Chapter 3

The Impact of Pricing and Buy-back Menus on Supply Chain Performance

3.1. Introduction

A variety of trade promotions are used regularly in supply chains to align the interests of retailers with the interests of the suppliers. Such trade deals provided by suppliers to retailers try to incentivize retailers to lower their prices and thereby increase sales. In many cases, this is done through promotions involving price discounts and other strategies that entail suppliers reducing list prices. Unfortunately, such price discount trade promotions are not viewed favorably by managers as they are seen to be "eroding the power of the brand" (see Ailawadi et al. [6]).

To better understand channel efficiency and other alternative incentive mechanisms in decentralized supply chains, we consider a single supplier multi-retailer supply chain selling a single product, in which retailers need to make pricing and inventory decisions. Our objective is to quantify the loss of efficiency due to decentralization, and to analyze the feasibility of incentive mechanisms, such as a *buy-back menu*, to improve supply chain efficiency.

We focus on a Stackelberg game wherein the supplier is the leader, and the retailers are followers that face stochastic demand, that is a function of their prices. In the simplest form

of this supply chain game, the supplier proposes a wholesale price to each of the retailers. The retailers then specify to the supplier their respective order quantities and decide their retail prices. Finally, demand for the period is realized. The supplier and the retailers are assumed to be selfish, rational, risk-neutral agents, seeking to maximize their expected profits.

A special case of our model is the single-supplier, single-retailer case when customer demand is not a function of retail price. It is widely known in the literature that *price-only contracts* do not coordinate the supply chain in this case [12, 70, 83, 85]. Lariviere and Porteus [70] and more recently, Perakis and Roels [85] present a detailed analysis of the supply chain efficiency in this game. On the other hand, Pasternack [83] showed that buy-back contracts *coordinate* the supply chain in this setting. A contract is said to *coordinate* the supply chain if the set of supply chain optimal actions is a Nash Equilibrium, i.e., no agent has a profitable unilateral deviation from the set of supply chain optimal actions (Cachon [12]).

The situation is different when demand is endogenous. Emmons and Gilbert [33] analyze the use of buy-back contracts ("returns policies") in an endogenous, *multiplicative* demand setting faced by a *single* retailer. Emmons and Gilbert demonstrate that under certain conditions, the performance of a buy-back contract is better than that of a price-only contract. However, they do not quantify theoretically the performance of a buy-back contract in that setting. Indeed, Bernstein and Federgruen [9] establish that in a single-supplier, single-retailer setting, a buy-back contract cannot coordinate the supply chain in any price-dependent demand scenario. However, Granot and Yin [48] evaluate the effectiveness of the buy-back contract in this two-echelon supply chain facing multiplicative demand. Particularly, they show that when the random part of the demand is uniformly distributed and the expected demand is a linear (resp., exponential) function of price, then the channel efficiency of the decentralized supply chain is 75% (resp., $\frac{2}{e} \approx 73.58\%$). Song et al. [95] further generalize the results of Granot and Yin [48] to the class of all multiplicative demand functions where the random part of the demand follows an *Increasing Generalized Failure*

100

Rate (IGFR) distribution. They also provide illustrative examples to show that none of the structural insights of buy-back contracts in multiplicative demand settings continue to hold for additive demand. In both these papers however, the authors leave open the analysis of the performance of a buy-back contract in an additive demand setting.

Bernstein and Federgruen [9] show that coordination is possible in a two-echelon supply chain facing additive (or multiplicative) price-dependent demand using a *price-discount sharing* (PDS) scheme coupled with a buy-back scheme. Specifically, the authors analyze a general model with a single supplier and multiple non-competing retailers and show that a linear PDS scheme along with a buy-back scheme may be used to coordinate the supply chain. Moreover, they show that even in the case of competing retailers, there exists a Nash Equilibrium for the retailers where coordination may again be achieved via a nonlinear PDS scheme coupled with a buy-back scheme.

As observed earlier, there are managerial limitations to the PDS scheme (Ailawadi et al. [6]). Granot and Yin [48] also note that, "the PDS scheme bears close resemblance to the traditional 'bill back' or 'count-recount' schemes, which, unfortunately are reported to be disliked by retailers (see, e.g., Blattberg and Neslin [10], Chapter 11)." Moreover, Ailawadi et al. [6] report that Everyday Low Purchase Price approaches (simple fixed wholesale price strategies) are widely preferred by managers for their simplicity. Addressing this fixed wholesale price approach, Bernstein and Federgruen [9] also identify conditions under which a fixed price scheme coupled with a buy-back scheme may coordinate the supply chain. However, in this case, "the constant pricing-scheme allows for only a single vector of wholesale prices" and this "may result in very small or zero margins for the supplier," see Bernstein and Federgruen [9]. Thus while the buy-back contract may be coordinating, it may yield very low margins of profit for the supplier, therefore deterring the supplier from adopting such an arrangement.

In this paper, we study the influence of a **buy-back menu** as an incentive in an *additive* demand setting with *multiple* coordinating or competing retailers. More specifically, we consider a payment scheme described by a constant per-unit wholesale price and a buy-back

menu, that is a *decreasing linear function of the retailer's selling price.* Intuitively, such a buy-back menu would tempt retailers to lower their prices (thus countering the effects of "double marginalization" (Spengler [96])) so that they may receive higher buy-back prices.

Our proposed incentive scheme has the advantage of retaining the virtues of an "every day low purchase price" (EDLPP) approach, that is viewed favorably by managers (Ailawadi et al. [6]). Moreover, the buy-back menu scheme does not have the problem of forward-buying, often associated with quantity-discount based trade deals, or the problem of "scam-backs" associated with price-discount based trade deals.

Unfortunately, the buy-back menu-based contract described above does not coordinate the supply chain. However in a linear demand setting, we establish worst-case bounds on the supply chain profit as well as on the supplier's profit in the Stackelberg game and show them to *outperform* price-only contracts. Specifically, we show that by using the suggested buy-back menu contract, the supplier's profit and therefore the supply chain profit is at least 50% of the optimal supply chain profit. Indeed, we show that no other contract involving a constant per unit wholesale price and a buy-back scheme can provide a better worst-case guarantee on the fraction of the global supply chain profit that the supplier can obtain. Furthermore, under a buy-back menu contract, we suggest heuristic wholesale pricing strategies for the supplier, using which the supplier may improve the performance of the entire supply chain to at least 75%, while still being guaranteed at least 50% of the optimal global supply chain profit. Thus, such a scheme may also provide for alternative sharing arrangements between supplier and retailer, subject to negotiations.

To quantify the performance of non-coordinating contracts in decentralized supply chains, we use the notion of *Price of Anarchy*, a term coined by Papadimitriou [82] and introduced in the theoretical computer science literature. Informally, the price of anarchy "measures the extent to which competition approximates cooperation" (Roughgarden [89]). More formally, the Price of Anarchy of a decentralized system is defined as the ratio of the optimal performance of a centralized system and the performance of a decentralized system at its worst Nash Equilibrium, with respect to the performance metric.

The Price of Anarchy concept was first studied by Koutsoupias and Papadimitriou [68] in the context of atomic games, wherein a common resource (a network) is shared by a finite number of players (wishing to send traffic across the network), each using an non-infinitesimal amount of it. Subsequently, this notion was extended to non-atomic games (Schmeidler [92]), wherein a resource is shared by an infinite number of users, each using an infinitesimal amount of it.

Some of the most important Price of Anarchy results to date have been for Congestion Games (Rosenthal [88]). See for example, Roughgarden and Tardos [90, 91], Correa et al. [23, 23, 24], and Perakis [84]. Other games for which the Price of Anarchy has been looked at include cost-sharing games (Moulin [77]), network resource allocation games (Johari and Tsitsiklis [59]), and network pricing games (Acemoglu and Ozdaglar [3]). In supply chain management literature, the use of the notion of price of anarchy has been relatively sparse (see (Martinez-de-Albéniz and Simchi-Levi [29]) and (Perakis and Roels [85])). Similar to Perakis and Roels [85], who use the Price of Anarchy as a measure to quantify the worst-case performance of price-only contracts in an exogenous demand setting, we use the Price of Anarchy concept to quantify the efficiency of a supply chain using a buy-back menu contract in an additive linear demand setting. Specifically, for a single retailer two-echelon supply chain using a buy-back menu contract, in the absence of explicit caps on retail price, the Price of Anarchy using a buy-back menu contract is $\frac{4}{3}$ (i.e., a relative efficiency of 75%). Thus the worst-case performance of a *buy-back menu contract* in an additive linear demand setting matches that of the *buy-back contract* in a multiplicative linear demand setting (Song et al. [95]). Furthermore, we show the same result also holds true for the case of multiple *coordinating* retailers. When *retailers compete*, we are able to establish that the Price of Anarchy of a supply chain using a buy-back menu is worse than that of a coordinating retailers system, but does not exceed 2 (i.e., a relative efficiency of 50%).

Finally we discuss the impact of vertical price constraints on supply chain performance in a buy-back menu regime. Specifically we study the Price of Anarchy of the supply chain configurations discussed earlier when retailers' maximum prices may be capped. We show

that price caps increase the price of anarchy in all the systems discussed to 2, while potentially eliminating any profit for the retailers.

3.2. Motivation

To motivate the loss of efficiency due to decentralization, we consider a simple model that though elementary, provides much of the governing intuition on the ramifications of decentralization. Consider a single product, single supplier, single retailer system facing *deterministic, endogenous* demand that is a linear function of the retail price, p. In other words, the retailer's demand is $d = a - bp$. Suppose that the supplier incurs a per unit cost of c for manufacturing the product. The supplier needs to decide on a wholesale price w at which she is willing to sell the product to the retailer, and subsequently, the retailer decides on his own retail price p, as well as inventory level, q. In this case, since the retailer possesses complete information regarding his demand once he has set his retail price, it follows that $q = d = a - bp$.

The Stackelberg game described above is easy to analyze and it follows that:

Observation 3.1. *In the system described above:*

- *The optimal retail price in the decentralized system, $\frac{3a+bc}{4b}$, is higher than the optimal retail price, $\frac{a+bc}{2b}$, in the centralized system*

- *The optimal order quantity in the decentralized system, $\frac{a-bc}{4}$, is only half of the optimal order quantity, $\frac{a-bc}{2}$, in the centralized system*

- *The optimal centralized system profit is $\frac{(a-bc)^2}{4b}$ and the optimal decentralized system profit is $\frac{3(a-bc)^2}{16b}$ which is 75% of the optimal centralized system profit*

- *The optimal expected supplier profit in the decentralized system is $\frac{(a-bc)^2}{8b}$ and this is 50% of the optimal centralized system profit.*

This supply chain efficiency loss due to decentralization is attributed to a phenomenon referred to in the literature as "double marginalization" (Spengler [96]). The question, of

course, is what is the impact of demand uncertainty on the loss of efficiency? Equally important, what is the effect of having many competing retailers on the loss of efficiency? To answer some of these questions, we make the following observation.

Observation 3.2. *There exists an instance with a single supplier and a single retailer facing stochastic demand, specified by $d = a - bp + \varepsilon$, in which the optimal decentralized expected system profit is only 18.75% of the optimal centralized expected system profit. Furthermore, the optimal supplier profit in the decentralized case is only 12.5% of the optimal centralized expected system profit.*

Given this huge drop in supply chain efficiency when demand is stochastic, we are motivated to study incentive contracts that improve the efficiency of the supply chain. Further, given the prevalence of an "Everyday Low Purchase Prices" approach in practice, we focus on the following question:

Are there "fixed wholesale price"- based incentive mechanisms that can help increase the efficiency of this decentralized supply chain?

In this work, we answer this question in the affirmative for a general single supplier multiple retailer model. Details of the model, our proposed scheme, and our results follow in the next section.

3.3. Preliminaries

We consider a supply chain with one supplier selling a single product to n retailers over a single period. Each retailer, say i, charges a price p_i for the product. We use the notation $\mathbf{x}$ to denote the vector $(x_1, x_2, \ldots, x_n)^t$, corresponding to any parameter x_i for retailer i. Accordingly, we denote the retail price vector by $\mathbf{p}$. In this work, the demand, $d_i(\mathbf{p})$, faced by retailer i is modeled as an linear function of $\mathbf{p}$, with demand uncertainty modeled as an additive random variable, ε_i. In other words, the demand vector $\mathbf{d}(\mathbf{p}) = \mathbf{a} - \mathbf{B}\mathbf{p} + \xi$, where $\mathbf{a}$ is a constant vector, $\mathbf{B} = [b_{ij}]_{n \times n}$ is a constant matrix and $\xi = (\varepsilon_1, \ldots, \varepsilon_n)$ is a vector of random variables. Without loss of generality, we assume that each ε_i is a *non-negative*

105

random variable with a continuous cumulative distribution function, F_i, and expected value, μ_i.

We study two scenarios in this work: one is which there are no bounds on the retailers prices, and one in which there is an upper bound p_i^u on each retailer's selling price, p_i. The motivation to consider retail price caps is due to a U.S. Supreme Court ruling in 1997 that overturned an earlier ruling banning manufacturers from setting retail price caps (Supreme Court [97], Felsenthal [35]), while maintaining the illegality of price floors. In the presence of price bounds, we assume that $\mathbf{a} - \mathbf{Bp} \geq \mathbf{0}$ for any $\mathbf{p} \leq \mathbf{p^u}$. Finally we assume that information regarding the demand structure, i.e. $\mathbf{a}$, $\mathbf{B}$, $\mathbf{p^u}$ and the distribution of ξ is symmetric, and therefore known to the supplier as well as the retailers.

In our model, any demand that cannot be satisfied from the inventory of the retailer is assumed to be lost. This is in contrast to the possibility of unfulfilled customers moving to a competing retailer with available inventory. Additionally, we assume that lost sales do not incur a penalty for the retailer. Excess inventory at retailer i may be salvaged at the end of the period at a per unit value of e_i.

Since ε_i is a non-negative random variable, the amount of inventory that retailer i must maintain for demand $d_i(\mathbf{p})$ is no less than the deterministic part of $d_i(\mathbf{p})$, $a_i - \sum_{1 \leq j \leq n} b_{ij} p_j$. In addition, let u_i denote the amount of inventory that retailer i maintains to meet the random part of it's demand, ε_i. The retailer's choice of u_i therefore determines its customer service level for the product.

We model the supply chain as a Stackelberg game in which the supplier is the *leader* and the retailers are the *followers*. The sequence of events is as follows. The supplier first offers a wholesale price w_i to retailer i and an incentive mechanism through a **buy-back menu** $s_i(w_i, \bullet)$. The buy-back menu offered to retailer i is a collection of per unit buy-back prices, each corresponding to a different retailer i selling price, paid by the supplier for each unit of inventory that retailer i salvages. After receiving the wholesale price w_i and the buy-back menu $s_i(w_i, \bullet)$, each retailer decides on its own selling price, p_i, and order quantity $a_i - \sum_{1 \leq j \leq n} b_{ij} p_j + u_i$. Furthermore, in the decentralized case, each retailer must submit

106

order quantities without knowing other retailers' selling prices and order quantities.

Upon receiving the order quantities, the supplier begins production and delivers items to retailer i at a per unit cost of $c_i(< p_i^u)$, for each i. This cost is borne by the supplier. Finally, demand for the period is realized and the retailers fulfill demand based on their available stocks. At the end of the period, the retailers send back any excess inventory to the supplier for a buy-back price $s_i(w_i, p_i)(\geq e_i)$ paid by the supplier to retailer i. The supplier salvages this inventory at a per-unit price of e_i for items received from retailer i.

We make the following assumptions regarding the matrix $\mathbf{B}$, that are commonly used in the literature [9, 74, 103, 104]:

Assumption 1 *(Substitutes)*: $b_{ii} > 0$ and $b_{ij} \leq 0$ for any $i, j = 1, 2, ..., n$ and $i \neq j$. That is, a retailer's price has a negative effect on its own demand, but a non-negative effect on other retailers' demands.

Assumption 2 *(Symmetry)*: $b_{ij} = b_{ji}$ for any $i, j = 1, 2, ..., n$ and $i \neq j$. That is, the cross-effects of retailers' prices on each other are symmetric. In fact, Vives [104](see also Vives [102]) shows that demand with symmetric cross-effects in an oligopoly setting arises as a natural consequence if the demand system is derived from the optimization problem of a representative consumer with a quasilinear utility function.

Assumption 3 *(Strict Diagonal Dominance)*: $\sum_{1 \leq j \leq n} b_{ij} > 0$ for $i = 1, 2, ..., n$. That is, a retailer's price has a higher effect on its own demand than the total effect of the prices of all other retailers. Analytically, the diagonal dominance condition allows for a contraction mapping method to show the uniqueness and stability of Nash Equilibria in a multi-retailer setting, see for e.g., Bernstein and Federgruen [9], Milgrom and Roberts [74], and Vives [103].

We now highlight a few important structural properties of the matrix $\mathbf{B}$. For notational convenience, we denote the set $\{1, 2, ..., n\}$ as N. Furthermore, for any subset $K \subseteq N$, we define the complement set as $\overline{K} = N \backslash K$. Consequently, for any vector $\mathbf{z}$, we may define $\mathbf{z}_K = (z_i)_{i \in K}$, and for any matrix $\mathbf{Z}$, $\mathbf{Z}_{K, \overline{K}} = (z_{ij})_{i \in K, j \in \overline{K}}$ and $\mathbf{Z}_K = (z_{ij})_{i, j \in K}$.

Observation 3.3. *Under assumptions 1,2, and 3, $\mathbf{B}$ is symmetric and positive definite. Consequently, $\mathbf{p}^t \mathbf{B} \mathbf{p}$ is a strictly convex function of $\mathbf{p}$.*

The above result implies that $\mathbf{B}$ is nonsingular, i.e. $\mathbf{B}^{-1}$ exists (Strang [98]). Furthermore, it also guarantees that $\mathbf{B}^{-1}$ is symmetric and positive definite. However, assumptions 1 and 3 allow us to make a stronger observation regarding $\mathbf{B}^{-1}$, that is not true in general for arbitrary SPD matrices.

Observation 3.4. *Under assumptions 1 and 3, $\mathbf{B}^{-1}$ exists and has non-negative entries. Additionally, all diagonal elements of $\mathbf{B}^{-1}$ are positive.*

Proof. Matrices that satisfy assumptions 1 and 3 in fact belong to class of *Leontief matrices* (see also M-matrices (Horn and Johnson [54])), which have the property that their inverse exists and has non-negative entries. This result for Leontief matrices is due to Samuelson's Substitution theorem (See for eg., chapters 7, 8, 9 of Koopmans [65], or Holley [53]).

To verify that the diagonal elements are positive, observe that the matrix corresponding to the cofactor of b_{ii} again satisfies assumptions 1 and 3 and so from the non-singularity of it, the positivity of b_{ii} follows. ∎

Note that since $\mathbf{B} = \mathbf{B}^t$, $\mathbf{B}^{-1}$ is also symmetric. Also, since $\mathbf{B}^{-1}$ has non-negative entries, it follows that if for any vector $\mathbf{x}$, $\mathbf{B}\mathbf{x} \geq \mathbf{a}$, then $\mathbf{x} \geq \mathbf{B}^{-1}\mathbf{a}$.

We study three different systems in this paper. The first is a *centralized system* in which retail prices $\mathbf{p}$ and inventory quantities $\mathbf{a} - \mathbf{B}\mathbf{p} + \mathbf{u}$ at the retailers are set by the system, so as to achieve the maximum possible expected system-wide (supplier plus retailers) profit. We henceforth refer to this optimal system-wide expected profit in the centralized system as *global optimal profit.*

The second is a *coordinated retailers system* in which the supplier sets the wholesale prices, $\mathbf{w}$, and a buy-back menu, $\mathbf{s}(\mathbf{w}, \bullet)$. The retailers then set the retailer prices $\mathbf{p}$ and inventories $\mathbf{a} - \mathbf{B}\mathbf{p} + \mathbf{u}$ together, so as to maximize the expected total retailer profit.

The third system of interest is a *competing retailers system*. In this decentralized system, after the supplier sets the wholesale prices $\mathbf{w}$ and the buy-back menu, $\mathbf{s}(\mathbf{w}, \bullet)$, each retailer i decides on its own retail price p_i and order quantity. Clearly, the question here is whether or

not there exists a Nash equilibrium, and how retailers can make ordering decisions without knowing each other's pricing strategy.

In the following sections we discuss each of the above systems. In Section 3.4, we characterize the global optimal policy for the centralized system. In Section 3.5, we introduce the buy-back menu and present our main results regarding the price of anarchy of the coordinated retailers system when a buy-back menu contract is employed. Further, we show that the price of anarchy of the coordinated retailers system can drop significantly in the presence of retail price caps. To remedy this situation, we develop a heuristic supplier wholesale policy for the coordinated system. Using this heuristic wholesale price policy, we establish that the expected profit of the supplier is no less than half of the global optimal profit, while that of the supplier and the retailers together is no less than three quarters of the global optimal profit.

In Section 3.6, we present our results regarding the price of anarchy of the competing retailers system when the supplier adopts a buy-back menu contract, in settings with and without retail price caps. Again, we suggest a heuristic supplier policy for the competing retailers system. We show that this heuristic allows the supplier to capture no less than half of the global optimal profit. Similarly, the competing retailers system expected profit is at least three quarters of the global optimal profit. Additionally, we also show that for a fixed supplier wholesale price, retailer prices drop when retailers compete, rather than coordinate. Finally, we show that the supplier can attain a higher expected profit with competing retailers than coordinated ones, while the customers continue to experience the same retail prices and service levels. We end with some concluding remarks in Section 3.7.

3.4.   The Centralized System

In a centralized system, with supplier costs, c, and salvage values, e, the system needs to decide on retail prices, p, and customer service levels, u, so as to maximize expected system profit. This problem is exactly the newsvendor problem with pricing, which has been studied extensively in the literature. For a detailed survey of the problem, the reader is directed to

Petruzzi and Dada [86].

In this setting, the system-wide expected profit for a given retail price $\mathbf{p}$ and inventory level $\mathbf{a} - \mathbf{Bp} + \mathbf{u}$ is given by:

$$\begin{aligned} \Pi^c(\mathbf{p}, \mathbf{u}) &= \mathbf{p}^t(\mathbf{a} - \mathbf{Bp} + \mathbf{u} - \mathbf{o}) - \mathbf{c}^t(\mathbf{a} - \mathbf{Bp} + \mathbf{u}) + \mathbf{e}^t\mathbf{o} \\ &= (\mathbf{p} - \mathbf{c})^t(\mathbf{a} - \mathbf{Bp} + \mathbf{u}) - (\mathbf{p} - \mathbf{e})^t\mathbf{o}, \end{aligned}$$

where $o_i \overset{def}{=} E[u_i - \varepsilon_i]^+$ denotes the expected overstock of retailer i. Additionally, let $v_i \overset{def}{=} \frac{E[u_i - \varepsilon_i]^+}{F_i(u_i)}$. The following observation relates $\mathbf{u}$, $\mathbf{v}$, and $\mathbf{o}$:

Observation 3.5. $\mathbf{u} \geq \mathbf{v} \geq \mathbf{o}$.

Proof. For each i, $u_i \geq u_i - E[\varepsilon_i | \varepsilon_i \leq u_i] = E[u_i | \varepsilon_i \leq u_i] - E[\varepsilon_i | \varepsilon_i \leq u_i] = \frac{E[u_i - \varepsilon_i]^+}{F_i(u_i)} = v_i$ and $v_i = \frac{E[u_i - \varepsilon_i]^+}{F_i(u_i)} \geq E[u_i - \varepsilon_i]^+ = o_i$. ∎

Let $\mathbf{p}^c$ and $\mathbf{u}^c$ be the vector of prices and inventory levels that maximize $\Pi^c(\mathbf{p}, \mathbf{u})$. Also, let $o_i^c = E[u_i^c - \varepsilon_i]^+$ and $v_i^c = \frac{E[u_i^c - \varepsilon_i]^+}{F_i(u_i^c)}$. We have the following result from the newsvendor with pricing model, that relates $\mathbf{p}^c$ and $\mathbf{u}^c$.

Lemma 3.6. (Whitin [106]) $F_i(u_i^c) = \frac{p_i^c - c_i}{p_i^c - e_i}$ *for* $i = 1, 2, ..., n$.

Proof. $\Pi^c(\mathbf{p}^c, \mathbf{u}) = (\mathbf{p}^c)^t(\mathbf{a} - \mathbf{Bp}^c + \mathbf{u} - \mathbf{o}) - \mathbf{c}^t(\mathbf{a} - \mathbf{Bp}^c + \mathbf{u}) + \mathbf{e}^t\mathbf{o}$

$= (\mathbf{p}^c - \mathbf{c})^t(\mathbf{a} - \mathbf{Bp}^c) + (\mathbf{p}^c - \mathbf{c})^t\mathbf{u} - (\mathbf{p}^c - \mathbf{e})^t\mathbf{o}$.

Since $\frac{\partial \Pi^c(\mathbf{p}^c, \mathbf{u})}{\partial u_i} = p_i^c - c_i - [p_i^c - e_i] F_i(u_i)$, is nonincreasing in u_i, $\Pi^c(\mathbf{p}^c, \mathbf{u})$ is concave in u_i and is at it's optimal value when $F_i(u_i) = \frac{p_i^c - c_i}{p_i^c - e_i}$. Hence $F_i(u_i^c) = \frac{p_i^c - c_i}{p_i^c - e_i}$ for $i = 1, 2, ..., n$. ∎

3.5. The Coordinated System

In this system, retailers set their price and inventory vectors *together* so as to maximize the expected total retailer profit. For a given wholesale price vector $\mathbf{w}$ and buy-back menu $\mathbf{s}(\mathbf{w}, \bullet)$

110

from the supplier, if the retailers set their retail prices at $\mathbf{p}$ and inventories at $\mathbf{a} - \mathbf{Bp} + \mathbf{u}$, their expected total profit is given by the expression:

$$\Pi^r \left(\mathbf{w}, \mathbf{s}(\mathbf{w}, \bullet), \mathbf{p}, \mathbf{u} \right) = \mathbf{p}^t (\mathbf{a} - \mathbf{Bp} + \mathbf{u} - \mathbf{o}) - \mathbf{w}^t (\mathbf{a} - \mathbf{Bp} + \mathbf{u}) + \mathbf{s} \left(\mathbf{w}, \mathbf{p} \right)^t \mathbf{o}$$

$$= (\mathbf{p} - \mathbf{w})^t (\mathbf{a} - \mathbf{Bp} + \mathbf{u}) - [\mathbf{p} - \mathbf{s} \left(\mathbf{w}, \mathbf{p} \right)]^t \mathbf{o}$$

In general, even for a constant buy-back menu $\mathbf{s}(\mathbf{w}, \bullet)$, i.e., when the buy-back payment is independent of the selling price, $\Pi^r \left(\mathbf{w}, \mathbf{s}(\mathbf{w}, \bullet), \mathbf{p}, \mathbf{u} \right)$ need not be jointly concave in $(\mathbf{p}, \mathbf{u})$. However, we propose a heuristic buy-back menu for the supplier that not only ensures concavity of the resulting retailer profit function, but also ensures that customers receive the same service level as they do in an optimal centralized system.

3.5.1 Proposed Buy-back Menu Contract

In Observations 1 and 2, we showed that randomness in demand can have large negative effects on the performance of a decentralized system. We now propose a buy-back menu that may be used by the supplier to counter these negative effects. Consider the buy-back menu specified by:

$$s_i^H \left(w_i, p_i \right) \overset{def}{=} p_i - \frac{p_i - w_i}{F_i \left(u_i^c \right)} \qquad \text{if } p_i \geq w_i \text{ for } i = 1, 2, ..., n$$

and w_i otherwise. For each i, $s_i^H \left(w_i, \bullet \right)$ is a nonincreasing function of p_i and encourages retailer i to set a lower retail price p_i and therefore, order a larger quantity of the product from the supplier. It is not hard to see that the buy-back value offered according to such a menu is always less than the wholesale price being offered by the supplier. A *buy-back menu contract* comprises a fixed per-unit wholesale price and a buy-back menu of prices for unsold inventory. Such a buy-back menu contract may be viewed as an alternative incentive mechanism between the simple buy-back contract (Pasternack [83]) and the Price-Discount Sharing (PDS) scheme coupled with buy-back (Bernstein and Federgruen [9]), in that it adopts the fixed per-unit wholesale price of a simple buy-back contract and a price-dependent

buy-back value that is also implicit in a PDS scheme.

Observe that the returns value specified by the buy-back menu $s_i^H(w_i, p_i)$ can be less than the salvage value, e_i, available in the market, especially when p_i is large. Ineffective buy-back prices such as these can be transformed into a buy-back price satisfying $s_i(w_i, p_i) \geq e_i$ as follows. Instead of charging the retailer a linear cost for returning the unused product, the supplier could announce a higher wholesale price and provide the end customers with a manufacturer "instant" rebate. Specifically, suppose w is the wholesale price the supplier wishes to select and the anticipated retailer price is p with $s_i^H(w, p) < e_i$. To transform the problem, the supplier selects some $\xi \geq e_i - s_i^H(w, p)$, announces a wholesale price of $w + \xi$ with buy-back menu $s_i(w + \xi, p_i) = s_i^H(w, p_i - \xi) + \xi$ for each feasible retail price $p_i - \xi$, and announces a manufacturer rebate of ξ. Note that when the retailer selects a retail price of $p_i + \xi$, the customer actually faces a product price of p_i per unit product due to the manufacturer rebate. The retailer obtains a profit of $p_i + \xi - (w + \xi) = p_i - w$ per unit product sold, but loses $w + \xi - s_i(w + \xi, p_i + \xi) = w - s_i^H(w_i, p_i)$ per unit of unused product. The supplier obtains a revenue of $(w + \xi) - \xi = w$ per unit product sold by the retailer, and a revenue of $(w + \xi) - s_i(w + \xi, p_i + \xi) = w - s_i^H(w_i, p_i)$ per unit product returned by the retailer. Thus, this scenario with a manufacturer rebate is equivalent to a situation in which the supplier sets a wholesale price of w with buy-back menu $s_i^H(w, \bullet)(< e_i)$, and the retailer sets his price as p_i, with no manufacturer rebate offered by the supplier.

We now proceed to show that for our choice of buy-back menu $\mathbf{s}^H(\mathbf{w}, \bullet)$, the optimal centralized system customer service level, $\mathbf{u}^c$, is the maximizer for the customer service level in a coordinated retailers system as well, for any feasible retail price vector $\mathbf{p}$. Hence, the retailers optimal decision regarding their customer service level is fixed as $\mathbf{u}^c$, and the retailers need only to decide on the price vector $\mathbf{p}$ that maximizes $\Pi^r(\mathbf{w}, \mathbf{p}) = \Pi^r(\mathbf{w}, \mathbf{s}^H(\mathbf{w}, \bullet), \mathbf{p}, \mathbf{u}^c)$, that we show is concave in $\mathbf{p}$.

Lemma 3.7. *For any* $\mathbf{p} \geq \mathbf{w}$, $\Pi^r(\mathbf{w}, \mathbf{s}^H(\mathbf{w}, \mathbf{p}), \mathbf{p}, \mathbf{u}) \leq \Pi^r(\mathbf{w}, \mathbf{s}^H(\mathbf{w}, \mathbf{p}), \mathbf{p}, \mathbf{u}^c)$.

Proof. $\Pi^r\left(\mathbf{w}, \mathbf{s}^H(\mathbf{w}, \mathbf{p}), \mathbf{p}, \mathbf{u}\right) = (\mathbf{p} - \mathbf{w})^t (\mathbf{a} - \mathbf{Bp} + \mathbf{u}) - \left[\mathbf{p} - \mathbf{s}^H(\mathbf{w}, \mathbf{p})\right]^t \mathbf{o}$. Since

$$\frac{\partial \Pi^r\left(\mathbf{w}, \mathbf{s}^H(\mathbf{w}, \mathbf{p}), \mathbf{p}, \mathbf{u}\right)}{\partial u_i} = p_i - w_i - \left[p_i - s_i^H(w_i, p_i)\right] F_i(u_i)$$

is nonincreasing in u_i, it follows that $\Pi^r\left(\mathbf{w}, \mathbf{s}^H(\mathbf{w}, \mathbf{p}), \mathbf{p}, \mathbf{u}\right)$ is concave in u_i and is maximized when $F_i(u_i) = \frac{p_i - w_i}{p_i - s_i^H(w_i, p_i)} = F_i(u_i^c)$. ∎

The reader would note that it is sufficient to consider the case when $\mathbf{p} \geq \mathbf{w}$ in order to determine the optimal retail price vector, because otherwise, those retailers with $p_i < w_i$ would have an incentive to all increase their prices to w_i, thereby improving their own profit as well as that of the system.

It is not hard to see that Lemma 3.7 in fact holds for any arbitrary additive demand function, $\mathbf{d}(\mathbf{p}) + \xi$. Moreover, for any feasible price vector $\mathbf{p}$ such that $\mathbf{d}(\mathbf{p}) \geq 0$, Lemma 3.7 also holds for any arbitrary multiplicative demand function, $[d_i(\mathbf{p}) \cdot \varepsilon_i]_{i=1}^n$. Thus, by offering the buy-back menu $\mathbf{s}^H(\mathbf{w}, \bullet)$, the supplier may ensure that it is optimal for the retailers to maintain the *same customer service levels* as they do in the centralized system. More intuitively, *the heuristic buy-back menu controls the effect of randomness in demand on the retailers' ordering decisions.* We now return our attention of the retailers' pricing decisions in a linear additive demand setting.

3.5.2 Evaluating the Optimal Retail Price Vector

By substituting the expression for $\mathbf{s}^H(\mathbf{w}, \mathbf{p})$, we may rewrite the expected retailer profit function as:

$$\begin{aligned}
\Pi^r(\mathbf{w}, \mathbf{p}) &= \Pi^r\left(\mathbf{w}, \mathbf{s}^H(\mathbf{w}, \bullet), \mathbf{p}, \mathbf{u}^c\right) \\
&= (\mathbf{p} - \mathbf{w})^t (\mathbf{a} - \mathbf{Bp} + \mathbf{u}^c) - \left[\mathbf{p} - \mathbf{s}^H(\mathbf{w}, \mathbf{p})\right]^t \mathbf{o}^c \\
&= (\mathbf{p} - \mathbf{w})^t (\mathbf{a} - \mathbf{Bp} + \mathbf{u}^c - \mathbf{v}^c).
\end{aligned}$$

Since $\mathbf{B}$ is SPD, it follows that $\Pi^r(\mathbf{w}, \mathbf{p})$ is strictly concave. The objective of the retailers

113

is now to maximize $\Pi^r(\mathbf{w}, \mathbf{p})$. Here, we must distinguish in the two circumstances that we analyze: in the absence and presence of caps on the retail prices. We present our results for these two cases in the following.

3.5.3 Absence of Price Caps

In the absence of upper bounds on the retailers' prices, solving for the optimal retail prices, $\mathbf{p}^*(\mathbf{w})$, that maximize the total expected profit of the retailers is an unconstrained concave quadratic maximization problem. We now present our main result regarding the performance of the coordinated retailers supply chain when the supplier offers a buy-back menu contract to the retailers:

Theorem 3.8. *In a coordinated retailers system where the supplier proposes a buy-back menu contract, and the retailers face linear additive stochastic demand, the total expected supply chain profit is at least 75% of the global optimal profit. Moreover, the total expected supplier profit is at least 50% of the global optimal profit.*

Proof. As discussed earlier, the centralized system profit, Π^c, is given by:

$$\Pi^c(\mathbf{p}^c, \mathbf{u}^c) = (\mathbf{p}^c - \mathbf{c})^t (\mathbf{a} - \mathbf{B}\mathbf{p}^c + \mathbf{u}^c) - (\mathbf{p}^c - \mathbf{e})^t \mathbf{o}^c$$

Since $\frac{o_i^c}{F_i(u_i^c)} = v_i^c$, we have that $o_i^c(p_i^c - e_i) = v_i^c(p_i^c - c_i)$. Consequently, we may rewrite Π^c as:

$$\Pi^c(\mathbf{p}^c, \mathbf{u}^c) = (\mathbf{p}^c - \mathbf{c})^t (\mathbf{a} - \mathbf{B}\mathbf{p}^c + \mathbf{u}^c - \mathbf{v}^c).$$

Furthermore,

$$\mathbf{p}^c = \frac{1}{2}\mathbf{B}^{-1}(\mathbf{a} + \mathbf{B}\mathbf{c} + \mathbf{u}^c - \mathbf{o}^c).$$

We shall now compare the global optimal profit with the expected system profit in the coordinated retailers system. For coordinated retailers in a buy-back menu regime, we have

that:

$$\Pi^r(\mathbf{w}, \mathbf{p}) = (\mathbf{p} - \mathbf{w})^t(\mathbf{a} - \mathbf{B}\mathbf{p} + \mathbf{u}^c - \mathbf{v}^c) \quad \text{and} \quad \mathbf{p}^*(\mathbf{w}) = \frac{1}{2}\mathbf{B}^{-1}(\mathbf{a} + \mathbf{B}\mathbf{w} + \mathbf{u}^c - \mathbf{v}^c)$$

where $\mathbf{p}^*(\mathbf{w})$ denotes the retail price vector that maximizes $\Pi^r(\mathbf{w}, \mathbf{p})$. Observe that for the retailers to maximize their profit, it must be the case that $\mathbf{p}^*(\mathbf{w}) \geq \mathbf{w}$. Indeed it is easy to verify that as long as $\mathbf{a} - \mathbf{B}\mathbf{w} \geq 0$, i.e. the supplier offers a wholesale price with nonnegative demand for each retailer, the retailers' expected profit function is well-behaved.

From Lemma 3.14, the supplier's expected profit in a coordinated retailers system is given by:

$$\Pi^s(\mathbf{w}) = (\mathbf{w} - \mathbf{c})^t(\mathbf{a} - \mathbf{B}\mathbf{p}(\mathbf{w}) + \mathbf{u}^c - \mathbf{v}^c) + (\mathbf{p}(\mathbf{w}) - \mathbf{p}^c)^t(\mathbf{v}^c - \mathbf{o}^c)$$

Substituting the expression for $\mathbf{p}^*(\mathbf{w})$ from above, it is easy to verify from the properties of $\mathbf{B}$ that $\Pi^s(\mathbf{w})$ is strictly concave in $\mathbf{w}$. Therefore, we have that:

$$\begin{aligned}
\nabla_{\mathbf{w}}\Pi^s(\mathbf{w}) &= \mathbf{a} - \mathbf{B}\mathbf{p}^*(\mathbf{w}) + \mathbf{u}^c - \mathbf{v}^c - \frac{1}{2}\mathbf{B}(\mathbf{w} - \mathbf{c}) + \frac{1}{2}(\mathbf{v}^c - \mathbf{o}^c) \\
&= \frac{\mathbf{a}}{2} - \mathbf{B}\mathbf{w} + \frac{\mathbf{u}^c}{2} - \frac{\mathbf{o}^c}{2} + \frac{1}{2}\mathbf{B}\mathbf{c}
\end{aligned}$$

Moreover, the wholesale price that maximizes the supplier's expected profit, $\Pi^s(\mathbf{w})$ is given by:

$$\mathbf{w}^* = \frac{1}{2}\mathbf{B}^{-1}(\mathbf{a} + \mathbf{B}\mathbf{c} + \mathbf{u}^c - \mathbf{o}^c) = \mathbf{p}^c$$

and correspondingly, the retailers' optimal decentralized retail price, $\mathbf{p}^*(\mathbf{w}^*)$ is given by:

$$\mathbf{p}^*(\mathbf{w}^*) = \frac{1}{2}\mathbf{B}^{-1}(\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c) + \frac{1}{2}\mathbf{p}^c \geq \mathbf{w}^* = \mathbf{p}^c$$

Since $\mathbf{p}^*(\mathbf{w}^*) \geq \mathbf{p}^c$ and $\mathbf{v}^c \geq \mathbf{o}^c$ (Observation 5), it follows from the expression for $\Pi^s(\mathbf{w}^*)$

115

that:

$$\begin{aligned}
\Pi^s(\mathbf{w}^*) &\geq (\mathbf{w}^* - \mathbf{c})^t(\mathbf{a} - \mathbf{B}\mathbf{p}^*(\mathbf{w}^*) + \mathbf{u}^c - \mathbf{v}^c) \\
&= (\mathbf{p}^c - \mathbf{c})^t(\mathbf{a} - \mathbf{B}\mathbf{p}^*(\mathbf{w}^*) + \mathbf{u}^c - \mathbf{v}^c) \\
&= (\mathbf{p}^c - \mathbf{c})^t(\mathbf{a} - \frac{1}{2}(\mathbf{a} + \mathbf{B}\mathbf{p}^c + \mathbf{u}^c - \mathbf{v}^c) + \mathbf{u}^c - \mathbf{v}^c) \\
&= \frac{1}{2}(\mathbf{p}^c - \mathbf{c})^t(\mathbf{a} - \mathbf{B}\mathbf{p}^c + \mathbf{u}^c - \mathbf{v}^c) = \frac{1}{2}\Pi^c(\mathbf{p}^c, \mathbf{u}^c)
\end{aligned}$$

In order to establish bounds on the system profit, we first rewrite the expression $(\mathbf{p}^*(\mathbf{w}^*) - \mathbf{c})$ as follows:

$$\begin{aligned}
(\mathbf{p}^*(\mathbf{w}^*) - \mathbf{c}) &= \frac{1}{2}\mathbf{B}^{-1}(\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c) + \frac{1}{2}\mathbf{p}^c - \mathbf{c} \\
&= \frac{1}{2}\mathbf{B}^{-1}(\mathbf{a} - \mathbf{B}\mathbf{c} + \mathbf{u}^c - \mathbf{v}^c) + \frac{1}{2}(\mathbf{p}^c - \mathbf{c}) \\
&= \frac{1}{2}\mathbf{B}^{-1}(\mathbf{a} + \mathbf{B}\mathbf{c} + \mathbf{u}^c - \mathbf{o}^c) - \mathbf{c} - \frac{1}{2}\mathbf{B}^{-1}(\mathbf{v}^c - \mathbf{o}^c) + \frac{1}{2}(\mathbf{p}^c - \mathbf{c}) \\
&= \frac{3}{2}(\mathbf{p}^c - \mathbf{c}) - \frac{1}{2}\mathbf{B}^{-1}(\mathbf{v}^c - \mathbf{o}^c) \qquad\qquad (3.1)
\end{aligned}$$

Now, we have that:

$$\begin{aligned}
&\Pi^r(\mathbf{w}^*, \mathbf{p}^*(\mathbf{w}^*)) + \Pi^s(\mathbf{w}^*) \\
=\ & (\mathbf{p}^*(\mathbf{w}^*) - \mathbf{c})^t(\mathbf{a} - \mathbf{B}\mathbf{p}^*(\mathbf{w}^*) + \mathbf{u}^c - \mathbf{v}^c) + (\mathbf{p}^*(\mathbf{w}^*) - \mathbf{p}^c)^t(\mathbf{v}^c - \mathbf{o}^c) \\
=\ & \frac{1}{2}(\mathbf{p}^*(\mathbf{w}^*) - \mathbf{c})^t(\mathbf{a} - \mathbf{B}\mathbf{p}^c + \mathbf{u}^c - \mathbf{v}^c) + (\mathbf{p}^*(\mathbf{w}^*) - \mathbf{p}^c)^t(\mathbf{v}^c - \mathbf{o}^c) \\
=\ & \frac{3}{4}(\mathbf{p}^c - \mathbf{c})^t(\mathbf{a} - \mathbf{B}\mathbf{p}^c + \mathbf{u}^c - \mathbf{v}^c) \qquad\qquad (3.2) \\
&- \frac{1}{4}(\mathbf{B}^{-1}(\mathbf{v}^c - \mathbf{o}^c))^t(\mathbf{a} - \mathbf{B}\mathbf{p}^c + \mathbf{u}^c - \mathbf{v}^c) + (\mathbf{p}^*(\mathbf{w}^*) - \mathbf{p}^c)^t(\mathbf{v}^c - \mathbf{o}^c) \\
=\ & \frac{3}{4}\Pi^c(\mathbf{p}^c, \mathbf{u}^c) + \frac{1}{4}(\mathbf{B}^{-1}(\mathbf{v}^c - \mathbf{o}^c))^t(\mathbf{a} - \mathbf{B}\mathbf{p}^c + \mathbf{u}^c - \mathbf{v}^c) \\
\geq\ & \frac{3}{4}\Pi^c(\mathbf{p}^c, \mathbf{u}^c)
\end{aligned}$$

where equation (3.2) follows by substituting equation (3.1) in the expression above, and the final inequality from Observations 4 and 5, and since $\mathbf{d}(\mathbf{p}^c) = \mathbf{a} - \mathbf{B}\mathbf{p}^c \geq \mathbf{0}$. ∎

116

Thus, we establish that the price of anarchy of a coordinated retailers system that uses a buy-back menu contract is $\frac{4}{3}$. Consider a special case of the coordinated retailers setting, where the number of retailers is 1. Stated alternately, the above result guarantees that if a supplier offers the proposed buy-back menu contract to a single retailer facing linear additive demand, and then acts as a leader in the Stackelberg game, then the supply chain profit in the decentralized system is at least 75% of the global optimal profit. Furthermore, the supplier is guaranteed at least 50% of the global optimal profit. Observe that in this case, the supplier selects it's wholesale price, w, selfishly. From the deterministic example considered in Observation 1, it also follows that these bounds are in fact tight. Thus, in a linear additive demand setting with a single retailer, the proposed *buy-back menu contract* matches the worst-case performance of a *buy-back contract* in a linear multiplicative demand setting (Song et al. [95]). However, while the optimal buy-back contract in a multiplicative setting was distribution-free, the parameters of our buy-back menu contract (a decreasing linear function of retail price) in an additive setting depend on the demand distribution. Nevertheless, our result is of interest since as observed by Granot and Yin [48] and Song et al. [95], the worst-case properties of the *buy-back contract* in a single retailer, multiplicative demand setting cease to hold in an additive demand setting even when the stochastic component of demand, ϵ, is uniformly distributed. Moreover, we establish our result about the performance of a buy-back menu contract even for multi-retailer coordinated systems.

3.5.4 Retail Price Caps

In a setting where the supplier imposes price caps, $\mathbf{p}^u \geq \mathbf{p}^c$, on the retailers prices, the objective of the retailers is to maximize $\Pi^r(\mathbf{w}, \mathbf{p})$ subject to these upper bound constraints on retail prices. It follows from the strict concavity of the objective function and the convexity of the underlying constraint set, that a local maximum would correspond to the unique global maximum for this problem.

To solve the retailers' problem, we begin by solving the unconstrained problem of maximizing $\Pi^r(\mathbf{w}, \mathbf{p})$. Since $\Pi^r(\mathbf{w}, \mathbf{p})$ is strictly concave in $\mathbf{p}$, it is maximized when $\mathbf{p} = \mathbf{p}^*(\mathbf{w}) =$

117

$\frac{1}{2}\left\{\mathbf{B}^{-1}(\mathbf{a}+\mathbf{Bw}+\mathbf{u}^c-\mathbf{v}^c)\right\}$. In the case that $\mathbf{p}^*(\mathbf{w}) < \mathbf{p}^u$, the optimal retail price vector for the constrained optimization problem is set to be the solution to the unconstrained optimization problem, i.e. $\mathbf{p}(\mathbf{w}) = \mathbf{p}^*(\mathbf{w})$.

Suppose instead that $\mathbf{p}^*(\mathbf{w}) \not< \mathbf{p}^u$. In this case, it follows from strict concavity that $\Pi^r(\mathbf{w},\mathbf{p})$ must attain its constrained maximum value at a local maximum point on the boundary $\mathbf{p}^*_{L(\mathbf{w})} = \mathbf{p}^u_{L(\mathbf{w})}$ for some $L(\mathbf{w}) \subseteq N$. Then by the Karush Kuhn Tucker (KKT) optimality conditions, which are both necessary and sufficient for the problem at hand, $\mathbf{p}_{\overline{L(\mathbf{w})}}$ must be set to $\mathbf{p}^*_{\overline{L(\mathbf{w})}}$, where $\mathbf{p}^*_{\overline{L(\mathbf{w})}}$ satisfies the condition that $\frac{\partial\Pi^r(\mathbf{w},\mathbf{p})}{\partial\mathbf{p}_{\overline{L(\mathbf{w})}}}\big|_{\mathbf{p}_{L(\mathbf{w})}=\mathbf{p}^u_{L(\mathbf{w})}} = 0$. Equivalently, from the strict concavity of $\Pi^r(\mathbf{w},\mathbf{p})$ in $\mathbf{p}_{\overline{L(\mathbf{w})}}$ after fixing $\mathbf{p}_{L(\mathbf{w})} = \mathbf{p}^u_{L(\mathbf{w})}$, we have that:

$$
\begin{aligned}
\mathbf{p}^*_{\overline{L(\mathbf{w})}} &= \frac{1}{2}\Big\{\mathbf{B}^{-1}_{\overline{L(\mathbf{w})}}\Big[\big(\mathbf{a}_{\overline{L(\mathbf{w})}}+\mathbf{u}^c_{\overline{L(\mathbf{w})}}-\mathbf{v}^c_{\overline{L(\mathbf{w})}}\big)\\
&\quad -\mathbf{B}_{\overline{L(\mathbf{w})},L(\mathbf{w})}\big(2\mathbf{p}^u_{L(\mathbf{w})}-\mathbf{w}_{L(\mathbf{w})}\big)\Big]+\mathbf{w}_{\overline{L(\mathbf{w})}}\Big\}\\
&< \mathbf{p}^u_{\overline{L(\mathbf{w})}}
\end{aligned}
\tag{3.3}
$$

Additionally, the KKT conditions also have that for each $i \in L(\mathbf{w})$ the partial derivative of $\Pi^r(\mathbf{w},\mathbf{p})$ with respect to p_i at $\mathbf{p}_{\overline{L(\mathbf{w})}} = \mathbf{p}^*_{\overline{L(\mathbf{w})}}$ and $\mathbf{p}_{L(\mathbf{w})\setminus\{i\}} = \mathbf{p}^u_{L(\mathbf{w})\setminus\{i\}}$ is no less than 0. Using a similar argument as above, this may be equivalently stated as:

$$
\begin{aligned}
p^*_i(\mathbf{w}) &= \frac{1}{2b_{ii}}\Big[a_i+u^c_i-v^c_i-\mathbf{B}_{\{i\},L(\mathbf{w})\setminus\{i\}}\big(2\mathbf{p}^u_{L(\mathbf{w})\setminus\{i\}}-\mathbf{w}_{L(\mathbf{w})\setminus\{i\}}\big)\\
&\quad -\mathbf{B}_{\{i\},\overline{L(\mathbf{w})}}\big(2\mathbf{p}^*_{\overline{L(\mathbf{w})}}-\mathbf{w}_{\overline{L(\mathbf{w})}}\big)+w_i\Big]\\
&\geq p^u_i
\end{aligned}
\tag{3.4}
$$

Since the above optimization problem is an instance of a convex quadratic programming problem, it can be solved in polynomial time using interior-point methods. To the best of our knowledge, the fastest interior-point algorithm to solve linear constrained convex quadratic programming is one that uses $O(n^3 L)$ arithmetic operations, where L is the number of bits needed to encode the input, see e.g., Monteiro and Adler [76] or Goldfarb and Liu [46].

However, the simplicity of the underlying constraints and the special structure of $\mathbf{B}$ allow us to propose a faster and simpler polynomial-time algorithm to determine a set $L(\mathbf{w})$ and $\mathbf{p}^*_{\overline{L(\mathbf{w})}}$ satisfying the KKT conditions. We divert our attention in the following subsection to the details of this algorithm.

A Polynomial-Time Algorithm to solve the Retailer's Problem

As noted earlier, our problem is a linear constrained convex quadratic programming problem, and hence if the unconstrained optimal solution is not feasible, then an optimal solution must lie on the boundary of the feasible region. We propose a simple "dual"-type active set approach, that turns out to solve the problem to optimality in polynomial time. Consider the following more general quadratic programming problem with only upper bound constraints on the decision variable, $\mathbf{p}$:

$$\text{maximize} \quad -\tfrac{1}{2}\mathbf{p}^t\mathbf{B}\mathbf{p} + \widetilde{\mathbf{a}}^t\mathbf{p} + \widetilde{\mathbf{c}} \qquad \text{(P)}$$
$$\text{subject to} \quad \mathbf{p} \leq \mathbf{p}^u$$

Additionally, suppose that in problem (P), $\widetilde{\mathbf{B}} = \tfrac{1}{2}(\mathbf{B} + \mathbf{B}^t)$ satisfies assumptions 1 and 3 (i.e., $\widetilde{\mathbf{B}}$ is a Leontief matrix). In this case, it is not hard to see that the problem (P) is a convex quadratic programming problem. The gradient of the objective function of (P) with respect to $\mathbf{p}$ is of the form $\widetilde{\mathbf{a}} - \widetilde{\mathbf{B}}\mathbf{p}$. We show that Algorithm A, described shortly, can be used to find the optimal boundary L^* and the optimal solution $\mathbf{p}^*$ of (P). Let Δ_i represent the n-dimensional unit vector in the i^{th} direction.

119

ALGORITHM A

(Initialization)

Step 0: Set $L_0 = \emptyset$ and $k = 0$.

Step 1: For $i = 1, 2, ..., n$, if $\Delta_i^t \widetilde{\mathbf{B}}^{-1} \widetilde{\mathbf{a}} \geq p_i^u$, then set $L_0 = L_0 \cup \{i\}$.

(While loop)

Step 2: Define $\mathbf{p}_{L_k}^k = \mathbf{p}_{L_k}^u$ and $\mathbf{p}_{\overline{L_k}}^k = \widetilde{\mathbf{B}}_{\overline{L_k}}^{-1}\left(\widetilde{\mathbf{a}}_{\overline{L_k}} - \widetilde{\mathbf{B}}_{\overline{L_k}, L_k}\mathbf{p}_{L_k}^u\right)$.

If $\frac{1}{b_{ii}}\left[\widetilde{a}_i - \widetilde{\mathbf{B}}_{\{i\}, L_k\backslash\{i\}}\mathbf{p}_{L_k\backslash\{i\}}^u - \widetilde{\mathbf{B}}_{\{i\}, \overline{L_k}}\mathbf{p}_{\overline{L_k}}^k\right] \geq \mathbf{p}_i^u \quad \forall i \in L_k$

or $L_k = \emptyset$ then goto **Step 3**.

Else set $L_{k+1} = L_k$, and for each $i \in L_k$, set $L_{k+1} = L_{k+1}\backslash\{i\}$ if

$\frac{1}{b_{ii}}\left[\widetilde{a}_i - \widetilde{\mathbf{B}}_{\{i\}, L_k\backslash\{i\}}\mathbf{p}_{L_k\backslash\{i\}}^u - \widetilde{\mathbf{B}}_{\{i\}, \overline{L_k}}\mathbf{p}_{\overline{L_k}}^k\right] < \mathbf{p}_i^u$

and repeat **Step 2** with $k = k + 1$.

(Termination)

Step 3: Set $L^* = L_k$ and $\mathbf{p}^* = \mathbf{p}^k$.

Informally, Algorithm A begins by guessing that the optimal boundary solution is at L_0. Based on this guess, it then computes the corresponding value of $\mathbf{p}_{\overline{L_0}}^0$ so as to satisfy the KKT conditions for $\overline{L_0}$. Having done so, the algorithm then proceeds to check whether the KKT conditions are satisfied for the indices in L_0, in a manner analogous to inequality (3.4). This is indicated in the first If-condition of **Step 2**. If indeed the KKT conditions are valid for indices in L_0, the algorithm terminates and exits to **Step 3**. If not, the algorithm updates the set of boundary indices L by removing the indices where the KKT condition is violated, and iterates.

It is clear that the algorithm terminates in at most n iterations, where $\mathbf{p}$ is a vector of dimension n. Also, it follows that the cost of each iteration is at most the order of a cost of inverting an $n \times n$ matrix, for which the current best known algorithm is $O(n^{2+\gamma})$ with $\gamma = 0.376$ (Coppersmith and Winograd [22]). Hence, the running time complexity of the Algorithm A is $O(n^{3+\gamma})$.

To prove correctness, all that is left is to ensure that the solution computed by the algorithm for $\mathbf{p}_{\overline{L^*}}$ is indeed an interior solution, analogous to inequality (3.3). Unsurprisingly,

the special structure of $\widetilde{\mathbf{B}}$ helps us to show just this.

Lemma 3.9. $\forall k$, $\mathbf{p}_{\overline{L_k}}^k < \mathbf{p}_{\overline{L_k}}^u$. *Consequently*, $\mathbf{p}_{\overline{L^*}}^* < \mathbf{p}_{\overline{L^*}}^u$.

The proof of Lemma 3.9 is via induction. For details, we refer the reader to Appendix A.0.2. As a consequence of this result, we may now state our main result regarding the algorithm:

Theorem 3.10. *Algorithm A determines in polynomial time the unique optimal solution to the convex quadratic programming problem (P).*

The uniqueness of the optimal solution for problem (P) and Theorem 3.10 imply that:

Corollary 3.11. *Starting with any L_0 with $\mathbf{p}_{\overline{L_0}}^0 < \mathbf{p}_{\overline{L_0}}^u$ instead of Step 1, Algorithm A will generate the unique boundary $L \subseteq N$ at which optimality conditions for (P) are satisfied.*

Returning to our main discussion, we may now easily solve the retailer's problem using Algorithm A. $\mathbf{p}$ is then set to be $\mathbf{p}(\mathbf{w})$ with $\mathbf{p}_{L(\mathbf{w})}(\mathbf{w}) = \mathbf{p}_{L(\mathbf{w})}^u$ and $\mathbf{p}_{\overline{L(\mathbf{w})}}(\mathbf{w}) = \mathbf{p}_{\overline{L(\mathbf{w})}}^*$. Furthermore, using the KKT conditions we may easily verify that for any feasible supplier wholesale price $\mathbf{w} \leq \mathbf{p}^u$, the corresponding retailer price satisfies $\mathbf{p}(\mathbf{w}) \geq \mathbf{p}$.

Lemma 3.12. *For any $\mathbf{w} \leq \mathbf{p}^u$, $\mathbf{p}(\mathbf{w}) \geq \mathbf{w}$.*

Proof. First, note that $\mathbf{w}_{L(\mathbf{w})} \leq \mathbf{p}_{L(\mathbf{w})}^u = \mathbf{p}_{L(\mathbf{w})}(\mathbf{w})$. By the KKT conditions, at $\mathbf{p} = \mathbf{p}(\mathbf{w})$, the partial derivative of $\Pi^r(\mathbf{w}, \mathbf{p})$ with respect to $\mathbf{p}_{\overline{L(\mathbf{w})}}$ is

$$\left[\mathbf{a}_{\overline{L(\mathbf{w})}} - \mathbf{B}_{\overline{L(\mathbf{w})},N} \mathbf{p}(\mathbf{w}) + \mathbf{u}_{\overline{L(\mathbf{w})}}^c - \mathbf{v}_{\overline{L(\mathbf{w})}}^c \right] - \mathbf{B}_{\overline{L(\mathbf{w})},N} \left[\mathbf{p}(\mathbf{w}) - \mathbf{w} \right] = 0. \qquad (3.5)$$

The feasibility of $\mathbf{p}(\mathbf{w})$ implies that $\mathbf{a}_{\overline{L(\mathbf{w})}} \geq \mathbf{B}_{\overline{L(\mathbf{w})},N} \mathbf{p}(\mathbf{w})$. Additionally, we have from Observation 3.5 that $\mathbf{u}_{\overline{L(\mathbf{w})}}^c - \mathbf{v}_{\overline{L(\mathbf{w})}}^c \geq 0$. Combining these facts with equation (3.5) and using Assumption 1 of $\mathbf{B}$, we have that:

$$\mathbf{B}_{\overline{L(\mathbf{w})}} \left[\mathbf{p}_{\overline{L(\mathbf{w})}}(\mathbf{w}) - \mathbf{w}_{\overline{L(\mathbf{w})}} \right] \geq \mathbf{B}_{\overline{L(\mathbf{w})},N} \left[\mathbf{p}(\mathbf{w}) - \mathbf{w} \right] \geq 0$$

and hence $\mathbf{p}_{\overline{L(\mathbf{w})}}(\mathbf{w}) - \mathbf{w}_{\overline{L(\mathbf{w})}} \geq 0$. ∎

Bounds on the Retailers' Profit

To develop bounds on the performance of the coordinated retailers system in a setting with price caps, we need to relate the sum of the supplier's and retailers' profits in the coordinated retailers system, to the centralized system profit. The following lemma serves as a first step in relating these quantities.

Given any index set $K \subseteq N$, we define $\Pi^r \left(\mathbf{c}, \mathbf{p}^K \right) \overset{def}{=} \max_{\mathbf{p}_{\overline{K}}} \{ \Pi^r \left(\mathbf{c}, \mathbf{p} \right) : \mathbf{p}_K = \mathbf{p}_K^c, \mathbf{p} \le \mathbf{p}^u \}$.

Lemma 3.13. *For any $K \subseteq N$, $\Pi^r \left(\mathbf{c} \right) \overset{def}{=} \max_{\mathbf{p} \le \mathbf{p}^u} \Pi^r \left(\mathbf{c}, \mathbf{p} \right) \ge \Pi^r \left(\mathbf{c}, \mathbf{p}^K \right) \ge \Pi^c \left(\mathbf{p}^c, \mathbf{u}^c \right)$ and*
$\mathbf{p} \left(\mathbf{c} \right) \le \mathbf{p}^K \le \mathbf{p}^c$.

The reader is referred to Appendix A.0.3 for a detailed proof of Lemma 3.13.

If a *centralized system* were facing a buy-back menu $\mathbf{s}^H (\mathbf{c}, \bullet)$ instead of a per unit salvage value, $\mathbf{e}$, then according to Lemma 3.13, the expected system profit would increase, i.e. $\Pi^r \left(\mathbf{c} \right) \ge \Pi^c \left(\mathbf{p}^c, \mathbf{u}^c \right)$. Alternately, if a supplier were to replace fixed salvage values by the proposed buy-back menu, then the expected retailer profit would increase, and additionally, the optimal retail prices would decrease, i.e. $\mathbf{p} \left(\mathbf{c} \right) \le \mathbf{p}^c$. This also supports our earlier claim that the proposed buy-back menu would encourage lower retail prices.

In Lemmas 3.7 and 3.13, we have seen how our proposed buy-back menu affects the retailers' decisions and the centralized system decisions, respectively, in the presence of demand randomness. In the following section, we shall try to "artificially" use the buy-back menu to control the effect of randomness on the *supplier's decisions* so that system efficiency is not sacrificed.

Supplier's Wholesale Price Decision

In the coordinated system, let the supplier's expected profit be denoted by $\Pi^s \left(\mathbf{w}, \mathbf{s} \left(\mathbf{w}, \bullet \right) \right)$, where $\mathbf{w}$ denotes the supplier's wholesale price and the buy-back menu is given by $\mathbf{s} \left(\mathbf{w}, \bullet \right)$. In using the proposed buy-back menu $\mathbf{s}^H \left(\mathbf{w}, \bullet \right)$, the supplier's expected profit may be written as:

$$\Pi^s \left(\mathbf{w} \right) \overset{def}{=} \Pi^s \left(\mathbf{w}, \mathbf{s}^H \left(\mathbf{w}, \bullet \right) \right) = \left(\mathbf{w} - \mathbf{c} \right)^t \left[\mathbf{a} - \mathbf{B} \mathbf{p} \left(\mathbf{w} \right) + \mathbf{u}^c \right] - \left[\mathbf{s}^H \left(\mathbf{w}, \mathbf{p} \left(\mathbf{w} \right) \right) - \mathbf{e} \right]^t \mathbf{o}^c.$$

To find the optimal wholesale prices for the supplier, we first consider optimizing the function

$$\widetilde{\Pi}^s\left(\mathbf{w}\right) \overset{def}{=} \left(\mathbf{w} - \mathbf{c}\right)^t \left[\mathbf{a} - \mathbf{B}\mathbf{p}\left(\mathbf{w}\right) + \mathbf{u}^c - \mathbf{v}^c\right].$$

We observe that $\widetilde{\Pi}^s\left(\mathbf{w}\right)$ is the expected profit of the supplier, if she were also offered a buy-back menu $\mathbf{s}^H(\mathbf{c}, \bullet)$, instead of a salvage value, $\mathbf{e}$, for unsold products. As we demonstrate below, if the supplier were to make decisions based on $\widetilde{\Pi}^s\left(\mathbf{w}\right)$, it would be able to "artificially" use the buy-back menu to also control the effect of demand randomness on its decision, that is, it would be able to motivate the retailers to order more than what they would order otherwise.

Lemma 3.14 aids us in doing precisely this, by establishing that for large enough retail prices, $\widetilde{\Pi}^s\left(\mathbf{w}\right)$ is a lower bound on $\Pi^s\left(\mathbf{w}\right)$.

Lemma 3.14. $\Pi^s\left(\mathbf{w}\right) = \widetilde{\Pi}^s\left(\mathbf{w}\right) + \left[\mathbf{p}\left(\mathbf{w}\right) - \mathbf{p}^c\right]^t \left(\mathbf{v}^c - \mathbf{o}^c\right)$. *Hence,* $\Pi^s\left(\mathbf{w}\right) \geq \widetilde{\Pi}^s\left(\mathbf{w}\right)$ *for* $\mathbf{p}\left(\mathbf{w}\right) \geq \mathbf{p}^c$ *and equality holds when* $\mathbf{p}\left(\mathbf{w}\right) = \mathbf{p}^c$.

Proof. Since $s_i^H\left(w_i, p_i\right) o_i^c = \left[p_i - \frac{p_i - w_i}{F_i\left(u_i^c\right)}\right] o_i^c = p_i o_i^c - \left(p_i - w_i\right) v_i^c$ for $i = 1, 2, ..., n$,

$$\Pi^s\left(\mathbf{w}\right) - \widetilde{\Pi}^s\left(\mathbf{w}\right) = \left(\mathbf{w} - \mathbf{c}\right)^t \mathbf{v}^c - \left[\mathbf{s}^H\left(\mathbf{w}, \mathbf{p}\left(\mathbf{w}\right)\right) - \mathbf{e}\right]^t \mathbf{o}^c$$

$$= \left(\mathbf{w} - \mathbf{c}\right)^t \mathbf{v}^c - \mathbf{p}\left(\mathbf{w}\right)^t \mathbf{o}^c + \left[\mathbf{p}\left(\mathbf{w}\right) - \mathbf{w}\right]^t \mathbf{v}^c + \mathbf{e}^t \mathbf{o}^c$$

$$= \left[\mathbf{p}\left(\mathbf{w}\right) - \mathbf{c}\right]^t \mathbf{v}^c - \left[\mathbf{p}\left(\mathbf{w}\right) - \mathbf{e}\right]^t \mathbf{o}^c$$

$$= \left[\mathbf{p}\left(\mathbf{w}\right) - \mathbf{p}^c\right]^t \left(\mathbf{v}^c - \mathbf{o}^c\right) + \left[\mathbf{p}^c - \mathbf{c}\right]^t \mathbf{v}^c - \left[\mathbf{p}^c - \mathbf{e}\right]^t \mathbf{o}^c$$

$$= \left[\mathbf{p}\left(\mathbf{w}\right) - \mathbf{p}^c\right]^t \left(\mathbf{v}^c - \mathbf{o}^c\right)$$

since $o_i^c = E\left[u_i^c - \varepsilon_i\right]^+ = v_i^c F_i\left(u_i^c\right) = v_i^c \frac{p_i^c - c_i}{p_i^c - e_i}$ for $i = 1, 2, ..., n$. ∎

Taken together, Lemmas 3.13 and 3.14 imply that the performance of a decentralized system with a buy-back menu $\mathbf{s}^H(\mathbf{c}, \bullet)$ also offered *to* the supplier is no better than a decentralized system with a salvage value $\mathbf{e}$ offered to the supplier, for large enough retail prices. A by-product of this observation is that the main results of the paper would still hold if the supplier were facing a buy-back menu $s_i(c_i, p_i) = p_i - \frac{p_i - c_i}{F_i}$ for some constant F_i, instead of $\mathbf{e}$. Note however, that in this case, the supplier would offer the retailers a buy-back menu of the form $s_i(w_i, p_i) = p_i - \frac{p_i - w_i}{F_i}$.

Different wholesale price vectors can induce the retailers to select the same retailer price

vector $\mathbf{p}$ when $p_i = p_i^u$ for some $i \in N$. In order to maximize the supplier's profit, we propose Lemma 3.15 to select the best wholesale price, $\widetilde{\mathbf{w}}$. The key idea of Lemma 3.15 is to raise the supplier's wholesale price to the maximum value, $\mathbf{p}^u$, for all those retailers whose retail price is at $\mathbf{p}^u$. To ensure that this change in the supplier's wholesale price does not change the retailer price vector, the wholesale price available to other retailers whose price is not at $\mathbf{p}^u$ needs to be changed appropriately. As we show shortly, this new supplier's price vector, $\widetilde{\mathbf{w}}$, is unilaterally the maximum wholesale price vector that induces the same retailer price vector.

Lemma 3.15. *For any* $\mathbf{w} \leq \mathbf{p}^u$, *set* $\widetilde{\mathbf{w}}_{L(\mathbf{w})} = \mathbf{p}_{L(\mathbf{w})}^u$ *and*

$$\widetilde{\mathbf{w}}_{\overline{L(\mathbf{w})}} = \mathbf{w}_{\overline{L(\mathbf{w})}} - \mathbf{B}_{\overline{L(\mathbf{w})}}^{-1} \mathbf{B}_{\overline{L(\mathbf{w})},L(\mathbf{w})} \left[\mathbf{p}_{L(\mathbf{w})}(\mathbf{w}) - \mathbf{w}_{L(\mathbf{w})} \right] \geq \mathbf{w}_{\overline{L(\mathbf{w})}}.$$

Then $\widetilde{\mathbf{w}} \leq \mathbf{p}^u$; *Moreover,* $\mathbf{p}(\widetilde{\mathbf{w}}) = \mathbf{p}(\mathbf{w})$.

Proof. Since $\mathbf{p}_{L(\mathbf{w})}(\mathbf{w}) = \mathbf{p}_{L(\mathbf{w})}^u$ and $\mathbf{p}_{\overline{L(\mathbf{w})}}(\mathbf{w}) < \mathbf{p}_{\overline{L(\mathbf{w})}}^u$, we have from the KKT optimality conditions for $\mathbf{r}_{\overline{L(\mathbf{w})}}(\mathbf{w})$ that:

$$
\begin{aligned}
0 &= \mathbf{a}_{\overline{L(\mathbf{w})}} + \mathbf{u}_{\overline{L(\mathbf{w})}}^c - \mathbf{v}_{\overline{L(\mathbf{w})}}^c - \mathbf{B}_{\overline{L(\mathbf{w})},N} \mathbf{p}(\mathbf{w}) - \mathbf{B}_{\overline{L(\mathbf{w})},N} \left[\mathbf{p}(\mathbf{w}) - \mathbf{w} \right] \\
&= \mathbf{a}_{\overline{L(\mathbf{w})}} + \mathbf{u}_{\overline{L(\mathbf{w})}}^c - \mathbf{v}_{\overline{L(\mathbf{w})}}^c - \mathbf{B}_{\overline{L(\mathbf{w})},N} \mathbf{p}(\mathbf{w}) \\
&\quad - \mathbf{B}_{\overline{L(\mathbf{w})}} \left[\mathbf{p}_{\overline{L(\mathbf{w})}}(\mathbf{w}) - \mathbf{w}_{\overline{L(\mathbf{w})}} \right] - \mathbf{B}_{\overline{L(\mathbf{w})},L(\mathbf{w})} \left[\mathbf{p}_{L(\mathbf{w})}(\mathbf{w}) - \mathbf{w}_{L(\mathbf{w})} \right] \\
&= \mathbf{a}_{\overline{L(\mathbf{w})}} + \mathbf{u}_{\overline{L(\mathbf{w})}}^c - \mathbf{v}_{\overline{L(\mathbf{w})}}^c - \mathbf{B}_{\overline{L(\mathbf{w})},N} \mathbf{p}(\mathbf{w}) \\
&\quad - \mathbf{B}_{\overline{L(\mathbf{w})}} \left\{ \mathbf{p}_{\overline{L(\mathbf{w})}}(\mathbf{w}) - \left(\mathbf{w}_{\overline{L(\mathbf{w})}} - \mathbf{B}_{\overline{L(\mathbf{w})}}^{-1} \mathbf{B}_{\overline{L(\mathbf{w})},L(\mathbf{w})} \left[\mathbf{p}_{L(\mathbf{w})}(\mathbf{w}) - \mathbf{w}_{L(\mathbf{w})} \right] \right) \right\} \\
&= \mathbf{a}_{\overline{L(\mathbf{w})}} + \mathbf{u}_{\overline{L(\mathbf{w})}}^c - \mathbf{v}_{\overline{L(\mathbf{w})}}^c - \mathbf{B}_{\overline{L(\mathbf{w})},N} \mathbf{p}(\mathbf{w}) - \mathbf{B}_{\overline{L(\mathbf{w})}} \left[\mathbf{p}_{\overline{L(\mathbf{w})}}(\mathbf{w}) - \widetilde{\mathbf{w}}_{\overline{L(\mathbf{w})}} \right] \quad (3.6) \\
&= \mathbf{a}_{\overline{L(\mathbf{w})}} + \mathbf{u}_{\overline{L(\mathbf{w})}}^c - \mathbf{v}_{\overline{L(\mathbf{w})}}^c - \mathbf{B}_{\overline{L(\mathbf{w})},N} \mathbf{p}(\mathbf{w}) - \mathbf{B}_{\overline{L(\mathbf{w})},N} \left[\mathbf{p}(\mathbf{w}) - \widetilde{\mathbf{w}} \right] \quad (3.7)
\end{aligned}
$$

where equation (3.7) follows by noting that $\mathbf{p}_{L(\mathbf{w})}(\mathbf{w}) = \mathbf{p}_{L(\mathbf{w})}^u = \widetilde{\mathbf{w}}_{L(\mathbf{w})}$. From the feasibility of $\mathbf{p}(\mathbf{w}) (\leq \mathbf{p}^u)$ and Observation 3.5, we have that $\mathbf{a}_{\overline{L(\mathbf{w})}} + \mathbf{u}_{\overline{L(\mathbf{w})}}^c - \mathbf{v}_{\overline{L(\mathbf{w})}}^c - \mathbf{B}_{\overline{L(\mathbf{w})},N} \mathbf{p}(\mathbf{w}) \geq 0$. Combining this fact with equation (3.6), it would imply that: $\mathbf{B}_{\overline{L(\mathbf{w})}} \left[\mathbf{p}_{\overline{L(\mathbf{w})}}(\mathbf{w}) - \widetilde{\mathbf{w}}_{\overline{L(\mathbf{w})}} \right] \geq 0$ and hence $\widetilde{\mathbf{w}}_{\overline{L(\mathbf{w})}} \leq \mathbf{p}_{\overline{L(\mathbf{w})}}(\mathbf{w}) \leq \mathbf{p}_{\overline{L(\mathbf{w})}}^u$. Now that we have

established that $\widetilde{\mathbf{w}} \leq \mathbf{p}^u$, it follows from Lemma 3.12 that $\widetilde{\mathbf{w}} \leq \mathbf{p}(\widetilde{\mathbf{w}})$. Hence it must be that $\mathbf{p}_{L(\mathbf{w})}(\widetilde{\mathbf{w}}) = \mathbf{p}_{L(\mathbf{w})}^u = \mathbf{p}_{L(\mathbf{w})}(\mathbf{w})$ and consequently, equation (3.7) provides the optimality condition for $\mathbf{p}_{\overline{L(\mathbf{w})}}(\widetilde{\mathbf{w}}) = \mathbf{p}_{\overline{L(\mathbf{w})}}(\mathbf{w})$. $\blacksquare$

The above result provides a procedure to determine the maximum supplier price vector for any feasible retailer price vector. Of course, the supplier needs also to determine what retail price vector to "push for", so as to optimize its own profit. However, this is a hard problem to solve for in closed form, since $L(\mathbf{w})$ is a set-valued function of $\mathbf{w}$. Moreover, we are able to show via an example that:

Theorem 3.16. *In the presence of price caps, the total expected supply chain profit of a coordinated retailers supply chain is at least 50% of the global optimal supply chain profit. Additionally, the supplier itself is guaranteed at least 50% of the global optimal supply chain profit. However, these bounds are tight.*

Proof. In the proof of Theorem 3.23, we construct a heuristic wholesale price strategy, using which the supplier is guaranteed at least 50% of the global optimal supply chain profit, thus implying the claim above. One would expect that if the supplier acts selfishly, rather than adopt a heuristic wholesale pricing strategy, then it would be able to obtain a larger portion of the supply chain profit. However, we show via an example that in fact the bounds claimed above are tight, meaning that in the worst-case the supplier can do no better than our heuristic wholesale pricing strategy and moreover, retailers may make no profit!

Consider an example of a deterministic single retailer system where the retailer's demand function is specified by $d(p) = 10 - p$, i.e., $a = 10$, $b = 1$. Let the supplier's cost of production be $c = 2$ and salvage value $e = 1$. Furthermore, let there be an upper bound on the retailer's price, given by $p^u = 6 + 2\sqrt{2} - \delta$ for some $0 \leq \delta < 0.8$. Since there is no stochasticity in demand, it follows that $u = 0$. In this case, the centralized system profit, $\Pi^c(p, u) = (p - c)(a - bp)$. Hence, $p^c = p(c) = \frac{1}{2b}(a + bc) = 6$, $u^c = 0$ and $\Pi^c(p^c, u^c) = \frac{1}{4b}(a - bc)^2 = 16$.

In a decentralized (coordinated retailers) system, the retailer's profit may be expressed as $\Pi^r(w, p) = (p - w)(a - bp)$, which attains its maximum value, $\frac{1}{4b}(a - bw)^2$, at $p = p^*(w) = \frac{1}{2b}(a + bw)$ if $\frac{1}{2b}(a + bw) \leq p^u$.

We now evaluate the supplier's optimal wholesale price decision in the Stackelberg game, when the supplier selects the wholesale price to maximize its own profit. We may split the supplier's optimal choice of w into two cases. When $\frac{1}{2b}(a + bw) < p^u$, the supplier's expected profit is

$$\Pi^s(w) = (w - c)(a - bp^*(w)) = \frac{1}{2}(w - c)(a - bw)$$

which attains its maximum value of $\frac{1}{8b}(a - bc)^2 = 8$ at $w = \frac{1}{2b}(a + bc) = 6$. However, when $\frac{1}{2b}(a + bw) \geq p^u$, as seen in Lemma 3.15, it is optimal for the supplier to set it's wholesale price at $w = p^u$, and therefore the supplier's maximum expected profit is $(p^u - c)(a - bp^u) = \left(4 + 2\sqrt{2} - \delta\right)\left(4 - 2\sqrt{2} + \delta\right) = 16 - \left(2\sqrt{2} - \delta\right)^2 = 8 + 4\sqrt{2}\delta - \delta^2$ if $p^u \geq 8$.

Hence, if $\delta \to 0^+$, then the optimal supplier wholesale price would be p^u, with $\Pi^s(w) = 8 = \frac{1}{2}\Pi^c(p^c, u^c)$ while $\Pi^r(w, p)$ approaches 0. ∎

Thus, the supplier's ability to set price caps on the prices charged by the retailers to the end customer increase the price of anarchy of the supply chain from $\frac{4}{3}$ to 2. Moreover, we see that price caps may significantly decrease the profits of the retailers. As an alternative to the supplier acting selfishly, we propose a feasible wholesale price $\mathbf{w}^{HC}$ that would guarantee the supplier 50% of the global supply chain profit, while also improving the total retailer expected profit. Note that $\mathbf{w}^{HC}$ need not optimize either $\Pi^s(\mathbf{w})$ or $\widetilde{\Pi}^s(\mathbf{w})$. However, as we have established via the example in Theorem 3.16, the supplier cannot do any better in terms of worst-case fraction of the global supply chain profit by selfishly optimizing its profit.

To determine $\mathbf{w}^{HC}$ and the corresponding retail price, $\mathbf{p}^{1/2}$ we consider a new construct whose optimal solution is the desired retail price vector. For any $0 \leq \rho \leq \frac{1}{2}$, suppose that the partial derivative of the supplier profit function w.r.t. to $\mathbf{p}$ is given by:

$$\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c - \mathbf{B}\mathbf{p} - (1 - \rho)\left[\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c - \mathbf{B}\mathbf{p}(\mathbf{c})\right] \quad . \tag{3.8}$$

We can apply Algorithm A to obtain the unique $L(\rho) \subseteq N$ with $\mathbf{p}^{\rho}_{L(\rho)} = \mathbf{p}^u_{L(\rho)}$ and

$$\mathbf{p}^{\rho}_{\overline{L(\rho)}} = \mathbf{B}^{-1}_{\overline{L(\rho)}}\left[\rho\left(\mathbf{a}_{\overline{L(\rho)}} + \mathbf{u}^c_{\overline{L(\rho)}} - \mathbf{v}^c_{\overline{L(\rho)}}\right) + (1 - \rho)\mathbf{B}_{\overline{L(\rho)},N}\mathbf{p}(\mathbf{c}) - \mathbf{B}_{\overline{L(\rho)},L(\rho)}\mathbf{p}^u_{L(\rho)}\right] < \mathbf{p}^u_{\overline{L(\rho)}}$$

126

so that the following KKT optimality conditions are satisfied:

$$\mathbf{a}_{\overline{L(\rho)}} + \mathbf{u}^c_{\overline{L(\rho)}} - \mathbf{v}^c_{\overline{L(\rho)}} - \mathbf{B}_{\overline{L(\rho)},N}\mathbf{p}^\rho - (1-\rho)\left[\mathbf{a}_{\overline{L(\rho)}} + \mathbf{u}^c_{\overline{L(\rho)}} - \mathbf{v}^c_{\overline{L(\rho)}} - \mathbf{B}_{\overline{L(\rho)},N}\mathbf{p}\,(\mathbf{c})\right] = 0 \quad (3.9)$$

and

$$\mathbf{a}_{L(\rho)} + \mathbf{u}^c_{L(\rho)} - \mathbf{v}^c_{L(\rho)} - \mathbf{B}_{L(\rho),N}\mathbf{p}^\rho - (1-\rho)\left[\mathbf{a}_{L(\rho)} + \mathbf{u}^c_{L(\rho)} - \mathbf{v}^c_{L(\rho)} - \mathbf{B}_{L(\rho),N}\mathbf{p}\,(\mathbf{c})\right] \geq 0 \quad (3.10)$$

An intuitive interpretation of equation (3.8) is as follows. In the deterministic model of Observation 1, we saw that whereas the optimal centralized order quantity was $\frac{a-bc}{2}$, the optimal decentralized order quantity was only $\frac{a-bc}{4}$, half of the optimal centralized order quantity. Motivated by this observation, in the analysis that follows, we select for the supplier the highest wholesale price, $\mathbf{w}$, in the coordinated system that targets an order quantity that is approximately $\frac{1}{2}$ of the optimal centralized order quantities. Note that this choice of $\mathbf{w}$ might not optimize either $\Pi^s(\mathbf{w})$ or $\widetilde{\Pi}^s(\mathbf{w})$. However, we show later that this choice still has the same worst-case guarantee as the best-possible choice of $\mathbf{w}$ for the supplier. To prove this result, we keep track of the retailer prices $\mathbf{p}^\rho$ in the coordinated system which target an order quantity that is approximately $(1-\rho)$ of the optimal centralized order quantity, and this is precisely the condition of equation (3.8).

Returning to the output of Algorithm A for the construct of equation (3.8), we now present a comparative statics result for $L(\rho)$ and $\mathbf{p}^\rho$:

Lemma 3.17. *For any* $0 \leq \rho < \widehat{\rho} \leq \frac{1}{2}$, $L(\rho) \subseteq L(\widehat{\rho})$ *and* $\mathbf{p}^\rho \leq \mathbf{p}^{\widehat{\rho}}$.

Proof. We have from the definition of $\mathbf{p}^{\widehat{\rho}}_{\overline{L(\widehat{\rho})}}$ that:

$$
\begin{aligned}
\mathbf{p}^u_{\overline{L(\widehat{\rho})}} \;>\; & \mathbf{p}^{\widehat{\rho}}_{\overline{L(\widehat{\rho})}} \\
=\; & \mathbf{B}^{-1}_{\overline{L(\widehat{\rho})}}\left\{\mathbf{a}_{\overline{L(\widehat{\rho})}} + \mathbf{u}^c_{\overline{L(\widehat{\rho})}} - \mathbf{v}^c_{\overline{L(\widehat{\rho})}} - (1-\widehat{\rho})\left[\mathbf{a}_{\overline{L(\widehat{\rho})}} + \mathbf{u}^c_{\overline{L(\widehat{\rho})}} - \mathbf{v}^c_{\overline{L(\widehat{\rho})}} - \mathbf{B}_{\overline{L(\widehat{\rho})},N}\mathbf{p}\,(\mathbf{c})\right]\right\} \\
& -\mathbf{B}^{-1}_{\overline{L(\widehat{\rho})}}\mathbf{B}_{\overline{L(\widehat{\rho})},L(\widehat{\rho})}\mathbf{p}^u_{L(\widehat{\rho})} \\
\geq\; & \mathbf{B}^{-1}_{\overline{L(\widehat{\rho})}}\left\{\mathbf{a}_{\overline{L(\widehat{\rho})}} + \mathbf{u}^c_{\overline{L(\widehat{\rho})}} - \mathbf{v}^c_{\overline{L(\widehat{\rho})}} - (1-\rho)\left[\mathbf{a}_{\overline{L(\widehat{\rho})}} + \mathbf{u}^c_{\overline{L(\widehat{\rho})}} - \mathbf{v}^c_{\overline{L(\widehat{\rho})}} - \mathbf{B}_{\overline{L(\widehat{\rho})},N}\mathbf{p}\,(\mathbf{c})\right]\right\} \\
& -\mathbf{B}^{-1}_{\overline{L(\widehat{\rho})}}\mathbf{B}_{\overline{L(\widehat{\rho})},L(\widehat{\rho})}\mathbf{p}^u_{L(\widehat{\rho})} \hspace{4cm} (3.11)
\end{aligned}
$$

where inequality (3.11) follows by observing that $\widehat{\rho} > \rho$ and $\mathbf{a} - \mathbf{B}\mathbf{p}(\mathbf{c}) > 0$. By Corollary 1, we may replace **Step 1** of Algorithm A with $L_0 = L(\widehat{\rho})$ to obtain $L(\rho)$, and it then follows that $L(\rho) \subseteq L_0 = L(\widehat{\rho})$. Consequently, $\mathbf{p}^{\widehat{\rho}}_{L(\widehat{\rho})} = \mathbf{p}^u_{L(\widehat{\rho})} \geq \mathbf{r}^{\rho}_{L(\widehat{\rho})}$.

Since $\overline{L(\widehat{\rho})} \subseteq \overline{L(\rho)}$, equation (3.9) implies that:

$$
\mathbf{a}_{\overline{L(\overline{\rho})}} + \mathbf{u}^c_{\overline{L(\overline{\rho})}} - \mathbf{v}^c_{\overline{L(\overline{\rho})}} - \mathbf{B}_{\overline{L(\overline{\rho})},N}\mathbf{p}^{\overline{\rho}} = (1-\overline{\rho})\left[\mathbf{a}_{\overline{L(\overline{\rho})}} + \mathbf{u}^c_{\overline{L(\overline{\rho})}} - \mathbf{v}^c_{\overline{L(\overline{\rho})}} - \mathbf{B}_{\overline{L(\overline{\rho})},N}\mathbf{p}\,(\mathbf{c})\right]
$$

for $\overline{\rho} = \rho, \widehat{\rho}$. Furthermore, since the RHS of the above equation is positive, we may conclude that:

$$
\begin{aligned}
0 \;\leq\; & \mathbf{B}_{\overline{L(\widehat{\rho})},N}\left(\mathbf{p}^{\widehat{\rho}} - \mathbf{p}^{\rho}\right) \\
=\; & \mathbf{B}_{\overline{L(\widehat{\rho})}}\left(\mathbf{p}^{\widehat{\rho}}_{\overline{L(\widehat{\rho})}} - \mathbf{p}^{\rho}_{\overline{L(\widehat{\rho})}}\right) + \mathbf{B}_{\overline{L(\widehat{\rho})},L(\widehat{\rho})}\left(\mathbf{p}^{\widehat{\rho}}_{L(\widehat{\rho})} - \mathbf{p}^{\rho}_{L(\widehat{\rho})}\right) \\
\leq\; & \mathbf{B}_{\overline{L(\widehat{\rho})}}\left(\mathbf{p}^{\widehat{\rho}}_{\overline{L(\widehat{\rho})}} - \mathbf{p}^{\rho}_{\overline{L(\widehat{\rho})}}\right)
\end{aligned}
$$

and hence $\mathbf{p}^{\widehat{\rho}}_{\overline{L(\widehat{\rho})}} - \mathbf{p}^{\rho}_{\overline{L(\widehat{\rho})}} \geq 0$. $\blacksquare$

In Observation 1, we also noted that when the decentralized system order quantity was $(1-\rho)$ of the optimal centralized order quantity, the decentralized system profit was $(1-\rho)(1+\rho)$ of the optimal centralized system profit. In Lemma 3.18 below, we show that even in the multi-retailer system with stochastic demand, if the supplier uses the proposed buy-back menu, and targets an order quantity that is $(1-\rho)$ of the optimal centralized order quantity, then the expected system-wide profit would be at least $(1-\rho)(1+\rho)$ of the optimal centralized system profit.

Let $0 = \rho_1 < \rho_2 < \rho_3 < \ldots \rho_{\theta-1} \leq \rho_\theta = \frac{1}{2}$ be the values at which $L(\rho)$ changes. That is, $L(\rho)$ is the same for $\rho_i \leq \rho < \rho_{i+1}$ for $i = 1, 2, \ldots, \theta-1$ but $L\left(\rho_i^-\right) \neq L(\rho_i)$ for $i = 2, 3, \ldots, \theta-1$. In addition, let $L(\rho_0) = \varnothing$ and $\widetilde{L}(\rho_i) = L(\rho_i) \backslash L(\rho_{i-1})$ for $i = 1, 2, \ldots, \theta-1$. Observe that $\mathbf{p}^{\rho_1} = \mathbf{p}(\mathbf{c})$.

Lemma 3.18. *For any* $\rho_i \leq \rho < \rho_{i+1}$ *for some* $i = 1, 2, \ldots, \theta-2$ *or* $\rho_{\theta-1} \leq \rho \leq \rho_\theta$,

$$(\mathbf{p}^\rho - \mathbf{c})\left[\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c - \mathbf{B}\mathbf{p}^\rho\right]^t$$
$$\geq \sum_{1 \leq j \leq i} (1 + \rho_j)(1 - \rho_j)\left[\mathbf{p}_{\widetilde{L}(\rho_j)}(\mathbf{c}) - \mathbf{c}_{\widetilde{L}(\rho_j)}\right]^t \left[\mathbf{a}_{\widetilde{L}(\rho_j)} + \mathbf{u}^c_{\widetilde{L}(\rho_j)} - \mathbf{v}^c_{\widetilde{L}(\rho_j)} - \mathbf{B}_{\widetilde{L}(\rho_j),N}\mathbf{p}(\mathbf{c})\right]$$
$$+ (1 + \rho)(1 - \rho)\left[\mathbf{p}_{\overline{L(\rho_i)}}(\mathbf{c}) - \mathbf{c}_{\overline{L(\rho_i)}}\right]^t \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N}\mathbf{p}(\mathbf{c})\right].$$

We refer the reader to Appendix A.0.4 for proof details.

With this structural insight into the properties of $\mathbf{p}^\rho$, we are now ready to answer questions regarding what retail price vector the supplier would like to induce via its wholesale price vector, and therefore regarding the underlying wholesale price decision of the supplier. Define a new wholesale price vector $\mathbf{w}^{1/2}$ as follows:

$$\mathbf{w}^{1/2}_{L(1/2)} = \mathbf{p}^u_{L(1/2)}$$
$$\mathbf{w}^{1/2}_{\overline{L(1/2)}} = 2\mathbf{p}^{1/2}_{\overline{L(1/2)}} - \mathbf{B}^{-1}_{\overline{L(1/2)}}\left[\mathbf{a}_{\overline{L(1/2)}} + \mathbf{u}^c_{\overline{L(1/2)}} - \mathbf{v}^c_{\overline{L(1/2)}} - \mathbf{B}_{\overline{L(1/2)},L(1/2)}\mathbf{p}^u_{L(1/2)}\right]$$

We show that the wholesale price vector defined above in terms of $\mathbf{p}^{1/2}$ is in fact one that induces a retail price of $\mathbf{p}^{1/2}$.

Lemma 3.19. $\mathbf{w}^{1/2} \leq \mathbf{p}^u$; *moreover,* $\mathbf{p}\left(\mathbf{w}^{1/2}\right) = \mathbf{p}^{1/2}$.

Proof. Observe that

$$0 = \mathbf{B}_{\overline{L(1/2)}}\left[\mathbf{w}^{1/2}_{\overline{L(1/2)}} - 2\mathbf{p}^{1/2}_{\overline{L(1/2)}} + \mathbf{B}^{-1}_{\overline{L(1/2)}}\left(\mathbf{a}_{\overline{L(1/2)}} + \mathbf{u}^c_{\overline{L(1/2)}} - \mathbf{v}^c_{\overline{L(1/2)}} - \mathbf{B}_{\overline{L(1/2)},L(1/2)}\mathbf{p}^u_{L(1/2)}\right)\right]$$
$$= \mathbf{a}_{\overline{L(1/2)}} + \mathbf{u}^c_{\overline{L(1/2)}} - \mathbf{v}^c_{\overline{L(1/2)}} - \mathbf{B}_{\overline{L(1/2)},N}\mathbf{p}^{1/2} - \mathbf{B}_{\overline{L(1/2)},N}\left(\mathbf{p}^{1/2} - \mathbf{w}^{1/2}\right)$$

The remainder of the proof follows in a similar manner to that of Lemma 3.15. ∎

It is important at this point to observe that $\mathbf{w}^{1/2}$ has been chosen so that the supplier may not be able to raise its wholesale price without affecting the retail price. Suppose now that $\mathbf{p}^{1/2} \geq \mathbf{p}^c$, i.e, our newly defined retail price $\mathbf{p}^{1/2}$ is at least as large as the retail price

in the centralized system for each retailer. We then show that it is sufficient for the supplier to set its wholesale price, $\mathbf{w}^{HC}$, to be $\mathbf{w}^{1/2}$.

Lemma 3.20. *If $\mathbf{p}^{1/2} \geq \mathbf{p}^c$, then the expected coordinated retailers system profit is at least 75% of the centralized system profit. Additionally, the supplier's expected profit in this case is at least 50% of the centralized system profit. More formally:*

$\Pi^s \left(\mathbf{w}^{HC} \right) + \Pi^r \left(\mathbf{w}^{HC} \right) \geq \frac{3}{4} \Pi^c \left(\mathbf{p}^c, \mathbf{u}^c \right)$ *and* $\Pi^s \left(\mathbf{w}^{HC} \right) \geq \frac{1}{2} \Pi^c \left(\mathbf{p}^c, \mathbf{u}^c \right)$.

Proof. If $\mathbf{p}^{1/2} \geq \mathbf{p}^c$, then we have from Lemma 3.14 that:

$$
\begin{aligned}
\Pi^s \left(\mathbf{w}^{HC} \right) + \Pi^r \left(\mathbf{w}^{HC} \right) &\geq \widetilde{\Pi}^s \left(\mathbf{w}^{HC} \right) + \Pi^r \left(\mathbf{w}^{HC} \right) \\
&= \left(\mathbf{p}^{1/2} - \mathbf{c} \right) \left[\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c - \mathbf{B}\mathbf{p}^{1/2} \right]
\end{aligned}
$$

Using Lemma 3.18, we are now able to prove the desired lower bound result as follows:

$$\left(\mathbf{p}^{1/2} - \mathbf{c} \right) \left[\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c - \mathbf{B}\mathbf{p}^{1/2} \right]$$

$$\geq \sum_{j=1}^{\theta} \left(1 + \rho_j \right) \left(1 - \rho_j \right) \left[\mathbf{p}_{\widetilde{L}(\rho_j)} \left(\mathbf{c} \right) - \mathbf{c}_{\widetilde{L}(\rho_j)} \right]^t \left[\mathbf{a}_{\widetilde{L}(\rho_j)} + \mathbf{u}^c_{\widetilde{L}(\rho_j)} - \mathbf{v}^c_{\widetilde{L}(\rho_j)} - \mathbf{B}_{\widetilde{L}(\rho_j),N}\mathbf{p} \left(\mathbf{c} \right) \right]$$

$$+ \left(1 + \tfrac{1}{2} \right) \left(1 - \tfrac{1}{2} \right) \left[\mathbf{p}_{\overline{L(1/2)}} \left(\mathbf{c} \right) - \mathbf{c}_{\overline{L(1/2)}} \right]^t \left[\mathbf{a}_{\overline{L(1/2)}} + \mathbf{u}^c_{\overline{L(1/2)}} - \mathbf{v}^c_{\overline{L(1/2)}} - \mathbf{B}_{\overline{L(1/2)},N}\mathbf{p} \left(\mathbf{c} \right) \right]$$

Furthermore since $0 \leq \rho_j \leq 1/2$, it follows that:

$$\geq \sum_{j=1}^{\theta} \frac{3}{4} \left[\mathbf{p}_{\widetilde{L}(\rho_j)} \left(\mathbf{c} \right) - \mathbf{c}_{\widetilde{L}(\rho_j)} \right]^t \left[\mathbf{a}_{\widetilde{L}(\rho_j)} + \mathbf{u}^c_{\widetilde{L}(\rho_j)} - \mathbf{v}^c_{\widetilde{L}(\rho_j)} - \mathbf{B}_{\widetilde{L}(\rho_j),N}\mathbf{p} \left(\mathbf{c} \right) \right]$$

$$+ \frac{3}{4} \left[\mathbf{p}_{\overline{L(1/2)}} \left(\mathbf{c} \right) - \mathbf{c}_{\overline{L(1/2)}} \right]^t \left[\mathbf{a}_{\overline{L(1/2)}} + \mathbf{u}^c_{\overline{L(1/2)}} - \mathbf{v}^c_{\overline{L(1/2)}} - \mathbf{B}_{\overline{L(1/2)},N}\mathbf{p} \left(\mathbf{c} \right) \right]$$

$$= \frac{3}{4} \left[\mathbf{p} \left(\mathbf{c} \right) - \mathbf{c} \right] \left[\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c - \mathbf{B}\mathbf{p} \left(\mathbf{c} \right) \right] \quad = \quad \frac{3}{4} \Pi^r \left(\mathbf{c} \right) \quad \geq \quad \frac{3}{4} \Pi^c \left(\mathbf{p}^c, \mathbf{u}^c \right).$$

where the last inequality follows from Lemma 3.13. It should be evident from the above analysis that the key to proving this result is the strong lower bound result provided by

130

Lemma 3.18.

We now proceed to provide a lower bound on the supplier's expected profit. When $\rho = \frac{1}{2}$, we may rewrite the optimality condition of $\mathbf{p}^\rho$ in equation (3.9) as:

$$\mathbf{a}_{\overline{L(1/2)}} + \mathbf{u}^c_{\overline{L(1/2)}} - \mathbf{v}^c_{\overline{L(1/2)}} - \mathbf{B}_{\overline{L(1/2)},N}\mathbf{p}^{1/2} = \frac{1}{2}\left[\mathbf{a}_{\overline{L(1/2)}} + \mathbf{u}^c_{\overline{L(1/2)}} - \mathbf{v}^c_{\overline{L(1/2)}} - \mathbf{B}_{\overline{L(1/2)},N}\mathbf{p}\left(\mathbf{c}\right)\right]$$

Substituting for $\mathbf{p}^{1/2}_{\overline{L(1/2)}}$ from the above equation, we have that:

$$
\begin{aligned}
\mathbf{w}^{HC}_{\overline{L(1/2)}} &= 2\mathbf{p}^{1/2}_{\overline{L(1/2)}} - \mathbf{B}^{-1}_{\overline{L(1/2)}}\left[\mathbf{a}_{\overline{L(1/2)}} + \mathbf{u}^c_{\overline{L(1/2)}} - \mathbf{v}^c_{\overline{L(1/2)}} - \mathbf{B}_{\overline{L(1/2)},L(1/2)}\mathbf{p}^u_{L(1/2)}\right] \\
&= \mathbf{p}_{\overline{L(1/2)}}\left(\mathbf{c}\right) + \mathbf{B}^{-1}_{\overline{L(1/2)}}\left(\mathbf{a}_{\overline{L(1/2)}} + \mathbf{u}^c_{\overline{L(1/2)}} - \mathbf{v}^c_{\overline{L(1/2)}}\right) \\
&\quad + \mathbf{B}^{-1}_{\overline{L(1/2)}}\mathbf{B}_{\overline{L(1/2)},L(1/2)}\mathbf{p}_{L(1/2)}\left(\mathbf{c}\right) - 2\mathbf{B}^{-1}_{\overline{L(1/2)}}\mathbf{B}_{\overline{L(1/2)},L(1/2)}\mathbf{p}^u_{L(1/2)} \\
&\quad - \mathbf{B}^{-1}_{\overline{L(1/2)}}\left[\mathbf{a}_{\overline{L(1/2)}} + \mathbf{u}^c_{\overline{L(1/2)}} - \mathbf{v}^c_{\overline{L(1/2)}} - \mathbf{B}_{\overline{L(1/2)},L(1/2)}\mathbf{p}^u_{L(1/2)}\right] \\
&= \mathbf{p}_{\overline{L(1/2)}}\left(\mathbf{c}\right) - \mathbf{B}^{-1}_{\overline{L(1/2)}}\mathbf{B}_{\overline{L(1/2)},L(1/2)}\left[\mathbf{p}^u_{L(1/2)} - \mathbf{p}_{L(1/2)}\left(\mathbf{c}\right)\right] \\
&\geq \mathbf{p}_{\overline{L(1/2)}}\left(\mathbf{c}\right) \tag{3.12}
\end{aligned}
$$

The last inequality follows from the non-negativity of $\mathbf{B}^{-1}$, assumption 1 regarding $\mathbf{B}$, and since $\mathbf{p}(\mathbf{c}) \leq \mathbf{p}^u$.

Hence $\Pi^s\left(\mathbf{w}^{HC}\right) \geq \widetilde{\Pi}^s\left(\mathbf{w}^{HC}\right) = \left(\mathbf{w}^{HC} - \mathbf{c}\right)^t\left[\mathbf{a} - \mathbf{B}\mathbf{p}\left(\mathbf{w}^{HC}\right) + \mathbf{u}^c - \mathbf{v}^c\right]$

Since we know that $\mathbf{p}(\mathbf{w}^{HC}) = \mathbf{p}^{1/2}$, the above expression can be split up as:

$$
\begin{aligned}
&= \left[\mathbf{p}^u_{L(1/2)} - \mathbf{c}_{L(1/2)}\right]^t\left[\mathbf{a}_{L(1/2)} + \mathbf{u}^c_{L(1/2)} - \mathbf{v}^c_{L(1/2)} - \mathbf{B}_{L(1/2),N}\mathbf{p}^{1/2}\right] \\
&\quad + \left[\mathbf{w}^{HC}_{\overline{L(1/2)}} - \mathbf{c}_{\overline{L(1/2)}}\right]^t\left[\mathbf{a}_{\overline{L(1/2)}} + \mathbf{u}^c_{\overline{L(1/2)}} - \mathbf{v}^c_{\overline{L(1/2)}} - \mathbf{B}_{\overline{L(1/2)},N}\mathbf{p}^{1/2}\right]
\end{aligned}
$$

Furthermore, equation (3.12) would lend that:

$$
\begin{aligned}
&\geq \left[\mathbf{p}^u_{L(1/2)} - \mathbf{c}_{L(1/2)}\right]^t\left[\mathbf{a}_{L(1/2)} + \mathbf{u}^c_{L(1/2)} - \mathbf{v}^c_{L(1/2)} - \mathbf{B}_{L(1/2),N}\mathbf{p}^{1/2}\right] \\
&\quad + \left[\mathbf{p}_{\overline{L(1/2)}}\left(\mathbf{c}\right) - \mathbf{c}_{\overline{L(1/2)}}\right]^t\left[\mathbf{a}_{\overline{L(1/2)}} + \mathbf{u}^c_{\overline{L(1/2)}} - \mathbf{v}^c_{\overline{L(1/2)}} - \mathbf{B}_{\overline{L(1/2)},N}\mathbf{p}^{1/2}\right]
\end{aligned}
$$

Using the optimality conditions of equations (3.9) and (3.10) for $\rho = 1/2$, we may now

131

replace $\mathbf{p}^{1/2}$ in favor of $\mathbf{p}(\mathbf{c})$ in the above inequality to obtain:

$$
\begin{aligned}
\Pi^s\left(\mathbf{w}^{HC}\right) \;\geq\; & \frac{1}{2}\left[\mathbf{p}^u_{L(1/2)} - \mathbf{c}_{L(1/2)}\right]^t \left[\mathbf{a}_{L(1/2)} + \mathbf{u}^c_{L(1/2)} - \mathbf{v}^c_{L(1/2)} - \mathbf{B}_{L(1/2),N}\mathbf{p}\left(\mathbf{c}\right)\right] \\
& +\frac{1}{2}\left[\mathbf{p}_{\overline{L(1/2)}}\left(\mathbf{c}\right) - \mathbf{c}_{\overline{L(1/2)}}\right]^t \left[\mathbf{a}_{\overline{L(1/2)}} + \mathbf{u}^c_{\overline{L(1/2)}} - \mathbf{v}^c_{\overline{L(1/2)}} - \mathbf{B}_{\overline{L(1/2)},N}\mathbf{p}\left(\mathbf{c}\right)\right] \\
\geq\; & \frac{1}{2}\left[\mathbf{p}_{L(1/2)}\left(\mathbf{c}\right) - \mathbf{c}_{L(1/2)}\right]^t \left[\mathbf{a}_{L(1/2)} + \mathbf{u}^c_{L(1/2)} - \mathbf{v}^c_{L(1/2)} - \mathbf{B}_{L(1/2),N}\mathbf{p}\left(\mathbf{c}\right)\right] \\
& +\frac{1}{2}\left[\mathbf{p}_{\overline{L(1/2)}}\left(\mathbf{c}\right) - \mathbf{c}_{\overline{L(1/2)}}\right]^t \left[\mathbf{a}_{\overline{L(1/2)}} + \mathbf{u}^c_{\overline{L(1/2)}} - \mathbf{v}^c_{\overline{L(1/2)}} - \mathbf{B}_{\overline{L(1/2)},N}\mathbf{p}\left(\mathbf{c}\right)\right] \\
=\; & \frac{1}{2}\Pi^r\left(\mathbf{c}\right) \;\geq\; \frac{1}{2}\Pi^c\left(\mathbf{p}^c,\mathbf{u}^c\right).
\end{aligned}
$$

$\blacksquare$

The lemma thus completely characterizes the worst-case performance for the supplier and the system when $\mathbf{p}^{1/2} \geq \mathbf{p}^c$. This case is relatively easy to analyze since, $\widetilde{\Pi}^s$ is a lower bound on the supplier's profit, Π^s.

This lower bound does not hold, however, if $\mathbf{p}^{1/2} \not\geq \mathbf{p}^c$ (recall Lemma 3.14), and so the above proof technique would fail in this case. Thus, when $\mathbf{p}^{1/2} \not\geq \mathbf{p}^c$, we suggest a new wholesale price, $\mathbf{w}^{HC}$ and a corresponding retail price, $\widetilde{\mathbf{p}}^{1/2} \geq \mathbf{p}^c$, that would in turn be favorable for the supplier.

Suppose now that $\mathbf{p}^{1/2} \not\geq \mathbf{p}^c$. Let G be the set of those indices where the above inequality holds, i.e., $G = \left\{i \in N : p_i^{1/2} < p_i^c\right\}$. Recall from Lemma 3.13 that we may now define a new retail price vector $\mathbf{p}^G \geq \mathbf{p}(\mathbf{c})$ with the property that $\mathbf{p}^G_G = \mathbf{p}^c_G$. Correspondingly, we modify our partial derivative construct of equation (3.8) to consider only those retail price vectors $\mathbf{p}$ such that $\mathbf{p}_G = \mathbf{p}^c_G$. In other words, we fix the retail prices of those retailers with indices in G to be $\mathbf{p}^c_G$. Let $\mathbf{a}^c_G = 0$ and $\mathbf{a}^c_{\overline{G}} = \mathbf{a}_{\overline{G}} - \mathbf{B}_{\overline{G},G}\mathbf{p}^c_G$.

To define a modified retail price vector $\widetilde{\mathbf{p}}^{1/2}$ that still leads to the retailers ordering approximately half of the centralized system order quantity, consider the following appropri-

ately modified partial derivative construct. For any $0 \leq \rho \leq \frac{1}{2}$, define:

$$\mathbf{a}_{\overline{G}}^c - \mathbf{u}_{\overline{G}}^c - \mathbf{v}_{\overline{G}}^c - \mathbf{B}_{\overline{G}}\mathbf{p}_{\overline{G}} - (1 - \rho)\left[\mathbf{a}_{\overline{G}}^c + \mathbf{u}_{\overline{G}}^c - \mathbf{v}_{\overline{G}}^c - \mathbf{B}_{\overline{G}}\mathbf{p}_{\overline{G}}^G\right]. \tag{3.13}$$

Intuitively, in the above construct we raise the retail prices of those retailers in the set G to be the same as it was in the centralized system case, i.e. $\mathbf{p}_G = \mathbf{p}_G^c$. However, since we have raised retail prices for the set G, we now need to determine the corresponding retail prices of the remaining retailers so their order quantity is approximately half of that in the centralized system. For any $\varsigma \subseteq \overline{G}$, let $\widetilde{\varsigma}$ represent those indices in $\overline{G}$ not present in ς, i.e. $\widetilde{\varsigma} = \overline{G} \backslash \varsigma \subseteq \overline{G}$. Again, we use Algorithm A to obtain the unique optimal boundary solution for the above partial derivative construct. That is, we determine the set of indices $\varsigma(\rho) \subseteq \overline{G}$ with $\widetilde{\mathbf{p}}_G^\rho = \mathbf{p}_G^c$, $\widetilde{\mathbf{p}}_{\varsigma(\rho)}^\rho = \mathbf{p}_{\varsigma(\rho)}^u$ and

$$\widetilde{\mathbf{p}}_{\widetilde{\varsigma(\rho)}}^\rho = \mathbf{B}_{\widetilde{\varsigma(\rho)}}^{-1}\left[\rho\left(\mathbf{a}_{\widetilde{\varsigma(\rho)}}^c + \mathbf{u}_{\widetilde{\varsigma(\rho)}}^c - \mathbf{v}_{\widetilde{\varsigma(\rho)}}^c\right) + (1 - \rho)\mathbf{B}_{\widetilde{\varsigma(\rho)},\overline{G}}\mathbf{p}_G^G - \mathbf{B}_{\widetilde{\varsigma(\rho)},\varsigma(\rho)}\mathbf{p}_{\varsigma(\rho)}^u\right] < \mathbf{p}_{\widetilde{\varsigma(\rho)}}^u$$

such that the following KKT conditions are satisfied:

$$\mathbf{a}_{\widetilde{\varsigma(\rho)}}^c + \mathbf{u}_{\widetilde{\varsigma(\rho)}}^c - \mathbf{v}_{\widetilde{\varsigma(\rho)}}^c - \mathbf{B}_{\widetilde{\varsigma(\rho)},\overline{G}}\widetilde{\mathbf{p}}_{\overline{G}}^\rho - (1 - \rho)\left[\mathbf{a}_{\widetilde{\varsigma(\rho)}}^c + \mathbf{u}_{\widetilde{\varsigma(\rho)}}^c - \mathbf{v}_{\widetilde{\varsigma(\rho)}}^c - \mathbf{B}_{\widetilde{\varsigma(\rho)},\overline{G}}\mathbf{p}_{\overline{G}}^G\right] = 0 \tag{3.14}$$

$$\mathbf{a}_{\varsigma(\rho)}^c + \mathbf{u}_{\varsigma(\rho)}^c - \mathbf{v}_{\varsigma(\rho)}^c - \mathbf{B}_{\varsigma(\rho),\overline{G}}\widetilde{\mathbf{p}}_{\overline{G}}^\rho - (1 - \rho)\left[\mathbf{a}_{\varsigma(\rho)}^c + \mathbf{u}_{\varsigma(\rho)}^c - \mathbf{v}_{\varsigma(\rho)}^c - \mathbf{B}_{\varsigma(\rho),\overline{G}}\mathbf{p}_{\overline{G}}^G\right] \geq 0 \tag{3.15}$$

We proceed to define a wholesale price vector $\mathbf{w}^{HC}$ that would indeed induce retailers to select their retail price to be $\widetilde{\mathbf{p}}^{1/2}$. Let $\tau = \left\{i \in N : \widetilde{p}_i^{1/2} = p_i^u\right\}$. Define $\mathbf{w}^{HC}$ as follows:

$\mathbf{w}_\tau^{HC} = \mathbf{p}_\tau^u$

$\mathbf{w}_{\widetilde{\tau}}^{HC} = 2\widetilde{\mathbf{p}}_{\widetilde{\tau}}^{1/2} - \mathbf{B}_{\widetilde{\tau}}^{-1}\left(\mathbf{a}_{\widetilde{\tau}} + \mathbf{u}_{\widetilde{\tau}}^c - \mathbf{v}_{\widetilde{\tau}}^c - \mathbf{B}_{\widetilde{\tau},\tau}\mathbf{p}_\tau^u\right).$

It is not hard to see, via a result similar to Lemma 3.19, that $\mathbf{p}\left(\mathbf{w}^{HC}\right) = \widetilde{\mathbf{p}}^{1/2}$. We shall now show that this new wholesale price satisfies our requirements regarding expected system and supplier profits. For this purpose, we need the following two lemmas.

Lemma 3.21. *If $\mathbf{p}^{1/2} \not\geq \mathbf{p}^c$, then $\widetilde{\mathbf{p}}^{1/2} \geq \mathbf{p}^{1/2}$.*

Lemma 3.21 verifies that $\widetilde{\mathbf{p}}^{1/2} \geq \mathbf{p}^c$, since $\widetilde{\mathbf{p}}_G^{1/2} = \mathbf{p}_G^c$ and $\widetilde{\mathbf{p}}_{\overline{G}}^{1/2} \geq \mathbf{p}_{\overline{G}}^{1/2} \geq \mathbf{p}_{\overline{G}}^c$. Consequently, we still have $\widetilde{\Pi}^s(\mathbf{w})$ is a lower bound on $\Pi^s(\mathbf{w})$. Details of the proof of Lemma 3.21 may be

found in Appendix A.0.5.

Lemma 3.22. *If* $\mathbf{p}^{1/2} \not\geq \mathbf{p}^c$, *then* $\mathbf{w}_G^{HC} - \mathbf{c}_G \geq \frac{1}{2}\left(\mathbf{p}_G^c - \mathbf{c}_G\right)$.

Recall that for those retailers in G, we were forced to set $\widetilde{\mathbf{p}}_G$ to be $\mathbf{p}_G^c$, so that it could not be tuned for retailers in G to order half of their centralized system inventory levels. However, Lemma 3.22 helps us in this matter by establishing that for our choice of $\mathbf{w}^{HC}$, the supplier's share of the profit margin of each item sold is at least half. A detailed proof of Lemma 3.22 may be found in Appendix A.0.6. We are now ready to state the main result of this section:

Theorem 3.23. $\Pi^s\left(\mathbf{w}^{HC}\right) + \Pi^r\left(\mathbf{w}^{HC}\right) \geq \frac{3}{4}\Pi^c\left(\mathbf{p}^c, \mathbf{u}^c\right)$ *and* $\Pi^s\left(\mathbf{w}^{HC}\right) \geq \frac{1}{2}\Pi^c\left(\mathbf{p}^c, \mathbf{u}^c\right)$.

Recall from Lemma 3.20 that when $\mathbf{p}^{1/2} \geq \mathbf{p}^c$, the supplier can guarantee at least 50% of the global optimization profit for itself, and 75% for the system by setting $\mathbf{w}^{HC} = \mathbf{w}^{1/2}$. Theorem 3.23 extends these results to the case when $\mathbf{p}^{1/2} \not\geq \mathbf{p}^c$. Indeed, if in this case the supplier selects $\mathbf{w}^{HC}$ to induce a modified retail price $\widetilde{\mathbf{p}}^{1/2}$, then the same system and supplier profit bounds can be established. The reader is directed to Appendix A.0.7 for a complete proof of this case.

Observation 1, pertaining to the deterministic demand scenario, implies that the results of Theorem 3.23 are tight. In fact, Observation 1 also suggests the supplier cannot obtain a better worst-case guarantee on its fraction of the global supply chain profit by optimizing its own profit, e.g., by optimally selecting $\mathbf{w}$, instead of the proposed $\mathbf{w}^{HC}$. In fact, as we establish in the example illustrated in Theorem 3.16, retailers may also benefit from the supplier's choice of $\mathbf{w}^{HC}$, rather than having the supplier maximizing its own profit.

3.6.   The Competing Retailers System

In the competing retailers system, each retailer is assumed to be a selfish, rational agent. Retailers set their own price and order quantity so as to maximize their own expected profits.

The expected profit of retailer i is given by:

$$\Pi_i^r\left(w_i, s_i\left(w_i, \bullet\right), \mathbf{p}, u_i\right) = p_i(a_i - \sum_{i=1}^n b_{ij}p_j + u_i - o_i) - w_i(a_i - \sum_{i=1}^n b_{ij}p_j + u_i) + s_i\left(w_i, p_i\right)o_i$$

$$= \left(p_i - w_i\right)\left(a_i - \sum_{i=1}^n b_{ij}p_j\right) + \left(p_i - w_i\right)u_i - \left[p_i - s_i\left(w_i, p_i\right)\right]o_i \quad (3.16)$$

Again, we consider the buy-back menu, $s_i^H\left(w_i, p_i\right) = p_i - \frac{p_i - w_i}{F_i\left(u_i^c\right)}$ for $i = 1, 2, ..., n$. Similar to Lemma 3.7, it is not hard to show that:

Lemma 3.24. *For* $\mathbf{p} \geq \mathbf{w}$, $\Pi_i^r\left(w_i, s_i^H\left(w_i, \bullet\right), \mathbf{p}, u_i\right) \leq \Pi_i^r\left(w_i, s_i^H\left(w_i, \bullet\right), \mathbf{p}, u_i^c\right)$.

Proof. If $\mathbf{p} \geq \mathbf{w}$, then $\frac{\partial \Pi_i^r\left(w_i, s_i^H\left(w_i, \bullet\right), \mathbf{p}, u_i\right)}{\partial u_i} = p_i - w_i - \left[p_i - s_i^H\left(w_i, p_i\right)\right]F_i\left(u_i\right)$ is decreasing in u_i and equals to zero when $F_i\left(u_i\right) = \frac{p_i - w_i}{p_i - s_i^H\left(w_i, p_i\right)} = F_i\left(u_i^c\right)$. ∎

As with Lemma 3.7, the results of Lemma 3.24 also continue to hold in arbitrary additive demand settings, $\mathbf{d}(\mathbf{p}) + \xi$, as well as arbitrary multiplicative demand settings, $\left[d_i(\mathbf{p}) \cdot \varepsilon_i\right]_{i=1}^n$, as long as $\mathbf{d}(\mathbf{p}) \geq 0$. In the following subsections, we analyze the performance of the buy-back menu contract in a competing retailers system with caps on retail prices. Most of the results of the case when there are no bounds on retail prices carry over from the analysis of the price caps case, and hence we omit discussing the case of no price caps in detail.

3.6.1 The Retail Price Equilibrium

By the definition of $s_i^H\left(w_i, p_i\right)$, it follows that:

$$\Pi_i^r\left(w_i, \mathbf{p}\right) = \Pi_i^r\left(w_i, s_i^H\left(w_i, \bullet\right), \mathbf{p}, u_i^c\right)$$

$$= \left(p_i - w_i\right)\left(a_i - \sum_{i=1}^n b_{ij}p_j\right) + \left(p_i - w_i\right)u_i^c - \left[p_i - s_i^H\left(w_i, p_i\right)\right]o_i^c$$

$$= \left(p_i - w_i\right)\left(a_i - \sum_{i=1}^n b_{ij}p_j\right) + \left(p_i - w_i\right)u_i^c - p_i o_i^c + p_i o_i^c - \left(p_i - w_i\right)v_i^c$$

$$= \left(p_i - w_i\right)\left(a_i - \sum_{i=1}^n b_{ij}p_j + u_i^c - v_i^c\right).$$

We have from the Debreu-Fan-Glicksberg theorem that since the retailers' strategy space $\mathbf{p} \leq \mathbf{p}^u$ is compact and convex, and from the concavity of retailers' payoffs that there exists at least one pure strategy Nash Equilibrium in this game (see for e.g., Fudenberg and Tirole [41]). Furthermore, the *Strong Diagonal Dominance* property of $\mathbf{B}$ implies that the best-response mapping corresponding to the retailers' profit functions is a contraction mapping on the retailers' entire strategy space, $\mathbf{p} \leq \mathbf{p}^u$ (see for e.g., Cachon and Netessine [13]). As a consequence, there exists a *unique* Nash Equilibrium for the retailers' game. Leveraging our assumptions regarding $\mathbf{B}$, we can even compute this Nash Equilibrium in polynomial time.

Let $\widehat{\mathbf{B}}$ be the diagonal matrix of $\mathbf{B}$. That is, $\widehat{b}_{ij} = 0$ for $i \neq j$ and $\widehat{b}_{ii} = b_{ii}$ for $i = 1, 2, ..., n$. Observe that in order for $\mathbf{p}^*$ to be a Nash Equilibrium, it must be that p_i^* is the best response corresponding to $\mathbf{p}_{-i}^*$. Considering the differential condition of the best response function, we have that the partial derivative of $\Pi_i^r(w_i, \mathbf{p})$ with respect to p_i for $i = 1, 2, ..., n$ may be expressed as:

$$\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c - \mathbf{B}\mathbf{p} - \widehat{\mathbf{B}}(\mathbf{p} - \mathbf{w}).$$

Equivalently, $\mathbf{p}^*$ must be a local maximum corresponding to an objective function whose partial derivative with respect to $\mathbf{p}$ is as indicated above, with the upper bound constraints $\mathbf{p} \leq \mathbf{p}^u$. It is not hard to see that $(\mathbf{B} + \widehat{\mathbf{B}})$ satisfies the three assumptions of $\mathbf{B}$. As a consequence, the equivalent optimization problem is a convex quadratic programming problem with upper bound constraints $\mathbf{p} \leq \mathbf{p}^u$. Hence our Algorithm A can be used to find the unique equilibrium price vector $\mathbf{p}^* = \mathbf{p}^e(\mathbf{w})$ and an index set $L^e(\mathbf{w}) \subseteq N$ such that

$$\mathbf{p}_{L^e(\mathbf{w})}^e(\mathbf{w}) = \mathbf{p}_{L^e(\mathbf{w})}^u$$

$$\mathbf{p}_{\overline{L^e(\mathbf{w})}}^e(\mathbf{w}) = \left(\mathbf{B}_{\overline{L^e(\mathbf{w})}} + \widehat{\mathbf{B}}_{\overline{L^e(\mathbf{w})}}\right)^{-1}\left(\mathbf{a}_{\overline{L^e(\mathbf{w})}} + \mathbf{u}_{\overline{L^e(\mathbf{w})}}^c - \mathbf{v}_{\overline{L^e(\mathbf{w})}}^c + \widehat{\mathbf{B}}_{\overline{L^e(\mathbf{w})},N}\mathbf{w}\right)$$

$$- \left(\mathbf{B}_{\overline{L^e(\mathbf{w})}} + \widehat{\mathbf{B}}_{\overline{L^e(\mathbf{w})}}\right)^{-1}\left(\mathbf{B}_{\overline{L^e(\mathbf{w})},L^e(\mathbf{w})} + \widehat{\mathbf{B}}_{\overline{L^e(\mathbf{w})},L^e(\mathbf{w})}\right)\mathbf{p}_{L^e(\mathbf{w})}^u < \mathbf{p}_{\overline{L^e(\mathbf{w})}}^u$$

By arguments discussed earlier, this equilibrium retail price $\mathbf{p}^e(\mathbf{w})$ satisfies the following optimality conditions :

$$\mathbf{a}_{L^e(\mathbf{w})} + \mathbf{u}_{L^e(\mathbf{w})}^c - \mathbf{v}_{L^e(\mathbf{w})}^c - \mathbf{B}_{L^e(\mathbf{w}),N}\mathbf{p}^e(\mathbf{w}) - \widehat{\mathbf{B}}_{L^e(\mathbf{w}),N}[\mathbf{p}^e(\mathbf{w}) - \mathbf{w}] \geq 0 \qquad (3.17)$$

and

$$a_{\overline{L^e(\mathbf{w})}} + u_{\overline{L^e(\mathbf{w})}}^c - v_{\overline{L^e(\mathbf{w})}}^c - B_{\overline{L^e(\mathbf{w})},N} \mathbf{p}^e(\mathbf{w}) - \widehat{B}_{\overline{L^e(\mathbf{w})},N} [\mathbf{p}^e(\mathbf{w}) - \mathbf{w}] = 0 \qquad (3.18)$$

To bound the system profit in the competing retailers case, we relate the *competing retailers* system profit with the *coordinated* system profit. Given wholesale price $\mathbf{w}$ in the coordinated system, we characterize a wholesale price $\widehat{\mathbf{w}}$ in the competing retailers system under which retailer prices are same in both settings. Note that the total expected profit of the retailers may be expressed as:

$$\sum_{i=1}^{n} \Pi_i^r(w_i, \mathbf{p}^e(\mathbf{w})) = [\mathbf{p}^e(\mathbf{w}) - \mathbf{w}]^t [\mathbf{a} - B\mathbf{p}^e(\mathbf{w}) + \mathbf{u}^c - \mathbf{v}^c] = \Pi_e^r(\mathbf{w}, \mathbf{p}^e(\mathbf{w})). \qquad (3.19)$$

Lemma 3.25. *For any* $\mathbf{w} \leq \mathbf{p}^u$, *let* $\widehat{\mathbf{w}}_{L(\mathbf{w})} = \mathbf{p}_{L(\mathbf{w})}^u$ *and*

$$\widehat{\mathbf{w}}_{\overline{L(\mathbf{w})}} = \mathbf{w}_{\overline{L(\mathbf{w})}} - \widehat{B}_{\overline{L(\mathbf{w})}}^{-1} \left(B_{\overline{L(\mathbf{w})}} - \widehat{B}_{\overline{L(\mathbf{w})}} \right) \left(\mathbf{p}_{\overline{L(\mathbf{w})}}(\mathbf{w}) - \mathbf{w}_{\overline{L(\mathbf{w})}} \right)$$
$$- \widehat{B}_{\overline{L(\mathbf{w})}}^{-1} B_{\overline{L(\mathbf{w})},L(\mathbf{w})} \left(\mathbf{p}_{L(\mathbf{w})}^u - \mathbf{w}_{L(\mathbf{w})} \right) \geq \mathbf{w}_{\overline{L(\mathbf{w})}},$$

then $\mathbf{p}^e(\widehat{\mathbf{w}}) = \mathbf{p}(\mathbf{w})$; *Moreover,* $\mathbf{p}^e(\widehat{\mathbf{w}}) \geq \widehat{\mathbf{w}}$.

The proof of Lemma 3.25 may be found in Appendix A.0.8.

We now verify the intuitive notion that $\mathbf{p}^e(\mathbf{w})$ is weakly monotonic in $\mathbf{w}$.

Lemma 3.26. *If* $\mathbf{w} \leq \mathbf{w}' \leq \mathbf{p}^u$, *then* $\mathbf{p}^e(\mathbf{w}) \leq \mathbf{p}^e(\mathbf{w}')$. *Consequently, for any* $\mathbf{w}$, $\mathbf{p}^e(\mathbf{w}) \leq \mathbf{p}^e(\widehat{\mathbf{w}}) = \mathbf{p}(\mathbf{w})$.

Lemma 3.26 implies that if the supplier fixes its wholesale price at $\mathbf{w}$, then retail prices drop when the retailers compete amongst themselves, i.e., customers prefer the competing retailers system over the coordinated retailer system. Lemma 3.26 may be easily verified using ideas similar to those in Lemma 3.17, and Corollary 3.11.

3.6.2 Supplier's Decision: Evaluating Wholesale Price

For a given wholesale price $\mathbf{w}$ and a buy-back menu $\mathbf{s}(\mathbf{w}, \bullet)$, let the supplier's expected profit be denoted by $\Pi_e^s(\mathbf{w}, \mathbf{s}(\mathbf{w}, \bullet))$. In using the proposed buy-back menu $\mathbf{s}^H(\mathbf{w}, \bullet)$, the

supplier's expected profit may be expressed as:

$$\Pi_e^s\left(\mathbf{w}\right) = \Pi_e^s\left(\mathbf{w}, \mathbf{s}^H\left(\mathbf{w}, \bullet\right)\right) = \left(\mathbf{w} - \mathbf{c}\right)^t\left[\mathbf{a} - \mathbf{B}\mathbf{p}^e\left(\mathbf{w}\right) + \mathbf{u}^c\right] - \left[\mathbf{s}^H\left(\mathbf{w}, \mathbf{p}^e\left(\mathbf{w}\right)\right) - \mathbf{e}\right]^t\mathbf{o}^c.$$

To obtain the heuristic wholesale prices, we first consider optimizing the function:

$$\widetilde{\Pi}_e^s\left(\mathbf{w}\right) = \left(\mathbf{w} - \mathbf{c}\right)^t\left[\mathbf{a} - \mathbf{B}\mathbf{p}^e\left(\mathbf{w}\right) + \mathbf{u}^c - \mathbf{v}^c\right], \tag{3.20}$$

which we show is a lower bound of $\Pi_e^s\left(\mathbf{w}\right)$ for large enough retailer prices $\mathbf{p}^e\left(\mathbf{w}\right)$.

Lemma 3.27. $\Pi_e^s\left(\mathbf{w}\right) = \widetilde{\Pi}_e^s\left(\mathbf{w}\right) + \left[\mathbf{p}^e\left(\mathbf{w}\right) - \mathbf{p}^c\right]^t\left(\mathbf{v}^c - \mathbf{o}^c\right).$ *Hence,* $\Pi_e^s\left(\mathbf{w}\right) \geq \widetilde{\Pi}_e^s\left(\mathbf{w}\right)$ *for* $\mathbf{p}^e\left(\mathbf{w}\right) \geq \mathbf{p}^c$ *and equality holds when* $\mathbf{p}^e\left(\mathbf{w}\right) = \mathbf{p}^c.$

Proof. Since $s_i^H\left(w_i, p_i\right)o_i^c = \left[p_i - \frac{p_i - w_i}{F_i\left(u_i^c\right)}\right]o_i^c = p_i o_i^c - \left(p_i - w_i\right)v_i^c$ for $i = 1, 2, ..., n$,

$\Pi_e^s\left(\mathbf{w}\right) - \widetilde{\Pi}_e^s\left(\mathbf{w}\right) = \left(\mathbf{w} - \mathbf{c}\right)^t\mathbf{v}^c - \left[\mathbf{s}^H\left(\mathbf{w}, \mathbf{p}^e\left(\mathbf{w}\right)\right) - \mathbf{e}\right]^t\mathbf{o}^c$

$= \left(\mathbf{w} - \mathbf{c}\right)^t\mathbf{v}^c - \mathbf{p}^e\left(\mathbf{w}\right)^t\mathbf{o}^c + \left[\mathbf{p}^e\left(\mathbf{w}\right) - \mathbf{w}\right]^t\mathbf{v}^c + \mathbf{e}^t\mathbf{o}^c$

$= \left[\mathbf{p}^e\left(\mathbf{w}\right) - \mathbf{c}\right]^t\mathbf{v}^c - \left[\mathbf{p}^e\left(\mathbf{w}\right) - \mathbf{e}\right]^t\mathbf{o}^c$

$= \left[\mathbf{p}^e\left(\mathbf{w}\right) - \mathbf{p}^c\right]^t\left(\mathbf{v}^c - \mathbf{o}^c\right) + \left[\mathbf{p}^c - \mathbf{c}\right]^t\mathbf{v}^c - \left[\mathbf{p}^c - \mathbf{e}\right]^t\mathbf{o}^c$

Since $o_i^c = E\left[u_i^c - \varepsilon_i\right]^+ = v_i^c F_i\left(u_i^c\right) = v_i^c \frac{p_i^c - c_i}{p_i^c - e_i}$ for $i = 1, 2, ..., n$,

$\left[\mathbf{p}^c - \mathbf{c}\right]^t\mathbf{v}^c - \left[\mathbf{p}^c - \mathbf{e}\right]^t\mathbf{o}^c = 0.$ ∎

The following Theorem summarizes the performance guarantees if the supplier uses the wholesale price vector $\mathbf{w}^{HD} = \widehat{\mathbf{w}^{HC}}$ (as defined in Lemma 3.25) and the proposed buy-back menu $\mathbf{s}^H\left(\mathbf{w}, \bullet\right).$

Theorem 3.28. *In the competing retailers system, we have that:*

$\Pi_e^s\left(\mathbf{w}^{HD}\right) + \sum_{i=1}^n \Pi_i^r\left(w_i^{HD}, \mathbf{p}^e\left(\mathbf{w}^{HD}\right)\right) = \Pi^s\left(\mathbf{w}^{HC}\right) + \Pi^r\left(\mathbf{w}^{HC}\right) \geq \frac{3}{4}\Pi^c\left(\mathbf{p}^c, \mathbf{u}^c\right)$ *and*
$\Pi_e^s\left(\mathbf{w}^{HD}\right) \geq \Pi^s\left(\mathbf{w}^{HC}\right) \geq \frac{1}{2}\Pi^c\left(\mathbf{p}^c, \mathbf{u}^c\right).$

Proof. We first note that $\mathbf{p}^e\left(\mathbf{w}^{HD}\right) = \mathbf{p}\left(\mathbf{w}^{HC}\right)$ by Lemma 3.25. Additionally, the Lemma also gives us that $\mathbf{w}^{HD} \geq \mathbf{w}^{HC}.$ By Theorem 3.23, we may assume w.l.o.g. that $\mathbf{p}(\mathbf{w}^{HC}) \geq \mathbf{p}^c.$ We set $\mathbf{u} = \mathbf{u}^c$ since we are using the proposed buy-back menu $\mathbf{s}^H\left(\mathbf{w}, \bullet\right)$ in the heuristic

schemes for the coordinated and competing retailers systems. In the competing retailers system, if the supplier offers a wholesale price of $\mathbf{w}^{HD}$, then it follows that the retailer order quantities are exactly the same as the retailer order quantities in a coordinated system with a wholesale price of $\mathbf{w}^{HC}$. So the total system profits in both systems are exactly the same, since transfer payments between supplier and retailers do not matter. This, together with Theorem 3.23, verifies the first claim of the theorem.

Then, from Lemma 3.27, and equations (3.19) and (3.20):

$$
\begin{aligned}
\Pi_e^s\left(\mathbf{w}^{HD}\right) + \sum_{i=1}^n \Pi_i^r\left(w_i^{HD}, \mathbf{p}^e\left(\mathbf{w}^{HD}\right)\right) \; &\geq \; \widetilde{\Pi}_e^s\left(\mathbf{w}^{HD}\right) + \sum_{i=1}^n \Pi_i^r\left(w_i^{HD}, \mathbf{p}^e\left(\mathbf{w}^{HD}\right)\right) \\
&= \; \Pi^c\left(\mathbf{p}^e\left(\mathbf{w}^{HD}\right), \mathbf{u}^c\right) \\
&= \; \Pi^c\left(\mathbf{p}\left(\mathbf{w}^{HC}\right), \mathbf{u}^c\right) \\
&= \; \widetilde{\Pi}^s\left(\mathbf{w}^{HC}\right) + \Pi^r\left(\mathbf{w}^{HC}\right) \\
&\geq \; \frac{3}{4}\Pi^c\left(\mathbf{p}^c, \mathbf{u}^c\right)
\end{aligned}
$$

where the last inequality may be observed from Theorem .

Since $\mathbf{w}^{HD} \geq \mathbf{w}^{HC}$ and $\mathbf{p}^e\left(\mathbf{w}^{HD}\right) = \mathbf{p}\left(\mathbf{w}^{HC}\right)$, it is easy to check that $\widetilde{\Pi}_e^s\left(\mathbf{w}^{HD}\right) \geq$ $\widetilde{\Pi}^s\left(\mathbf{w}^{HC}\right)$. Also, since $\mathbf{p}(\mathbf{w}^{HC}) \geq \mathbf{p}^c$, we have that $\Pi_e^s\left(\mathbf{w}^{HD}\right) \geq \widetilde{\Pi}_e^s\left(\mathbf{w}^{HD}\right)$. Putting these together:

$$
\Pi_e^s\left(\mathbf{w}^{HD}\right) \geq \widetilde{\Pi}_e^s\left(\mathbf{w}^{HD}\right) \geq \widetilde{\Pi}^s\left(\mathbf{w}^{HC}\right) \geq \frac{1}{2}\Pi^c\left(\mathbf{p}^c, \mathbf{u}^c\right)
$$

where the last two inequalities may be seen from Theorem 3.23. ∎

Similar to the coordinated system case, Observation 1 also implies that there exist instances of the competing retailers system for which the bounds in Theorem 3.28 are tight.

We note some interesting implications of Theorem 3.23 and Theorem 3.28. In both the coordinated and competing retailers systems, the proposed heuristic for the supplier's wholesale price need not optimize its profit, $\Pi_e^s(\mathbf{w})$. Hence, while the supplier can perhaps do better for itself by selecting an alternate $\mathbf{w}^*$ that optimizes it's profit, this need not be

139

beneficial for either the supply chain or the retailers.

Indeed, the two theorems imply that if the supplier optimizes its wholesale price, she can guarantee *at least half of the global optimization profit*. However, this may lead to a supply chain expected profit arbitrarily close to the same amount, thus eliminating any profit for the retailers. We illustrate this through the example below.

Example 1: Consider deterministic single retailer system where the retailer's demand function is specified by $d(p) = 10 - p$, i.e., $a = 10$, $b = 1$. Let the supplier's cost of production be $c = 2$ and salvage value $e = 1$. Furthermore, let there be an upper bound on the retailer's price, given by $p^u = 6 + 2\sqrt{2} - \delta$ for some $0 \le \delta < 0.8$. Since there is no stochasticity in demand, it follows that $u = 0$. In this case, the centralized system profit, $\Pi^c(p, u) = (p - c)(a - bp)$. Hence, $p^c = p(c) = \frac{1}{2b}(a + bc) = 6$, $u^c = 0$ and $\Pi^c(p^c, u^c) = \frac{1}{4b}(a - bc)^2 = 16$.

In a decentralized (coordinated retailers) system, the retailer's profit may be expressed as $\Pi^r(w, p) = (p - w)(a - bp)$, which attains its maximum value, $\frac{1}{4b}(a - bw)^2$, at $p = p^*(w) = \frac{1}{2b}(a + bw)$ if $\frac{1}{2b}(a + bw) \le p^u$.

To determine the proposed heuristic retail price, $p^{1/2}$, we must satisfy the condition that $a - bp = \frac{1}{2}[a - bp(c)]$. This condition would imply that $p = \frac{a}{2b} + \frac{1}{2}p(c) = 5 + 3 = 8 < p^u$. Hence $p^{1/2} = 8$. Since $p^{1/2} > p^c$, the heuristic wholesale price would then be specified by $w^{HC} = w^{HD} = w^{1/2} = 2p^{1/2} - \frac{a}{b} = 16 - 10 = 6$. Hence the supplier's profit in the decentralized system is $\Pi^s(w^{HC}) = \Pi^s_e(w^{HD}) = (w^{1/2} - c)(a - bp^{1/2}) = (6 - 2)(10 - 8) = 8 = \frac{1}{2}\Pi^c(p^c, u^c)$. In addition, the retailer's profit in a decentralized system would be $\Pi^r(w^{HC}) = \Pi^r_1(w^{HD}, p^e(w^{HD})) = (p^{1/2} - w^{1/2})(a - bp^{1/2}) = (8 - 6)(10 - 8) = 4$. Hence, the total system profit in this case is $\Pi^s(w^{HC}) + \Pi^r(w^{HC}) = \Pi^s_e(w^{HD}) + \Pi^r_1(w^{HD}, p^e(w^{HD})) = 8 + 4 = 12 = \frac{3}{4}\Pi^c(p^c, u^c)$.

We now evaluate the supplier's optimal wholesale price decision in the Stackelberg game, when the supplier selects the wholesale price to maximize its own profit. We may split the supplier's optimal choice of w into two cases. When $\frac{1}{2b}(a + bw) < p^u$, the supplier's expected

profit is

$$\Pi^s(w) = (w - c)(a - bp^*(w)) = \frac{1}{2}(w - c)(a - bw)$$

which attains its maximum value of $\frac{1}{8b}(a - bc)^2 = 8$ at $w = \frac{1}{2b}(a + bc) = 6$. However, when $\frac{1}{2b}(a + bw) \geq p^u$, as seen in Lemma 3.15, it is optimal for the supplier to set it's wholesale price at $w = p^u$, and therefore the supplier's maximum expected profit is $(p^u - c)(a - bp^u) = (4 + 2\sqrt{2} - \delta)(4 - 2\sqrt{2} + \delta) = 16 - (2\sqrt{2} - \delta)^2 = 8 + 4\sqrt{2}\delta - \delta^2$ if $p^u \geq 8$.

Hence, if $\delta \rightarrow 0^+$, then the optimal supplier wholesale price would be p^u, with $\Pi^s(w) = 8 = \frac{1}{2}\Pi^c(p^c, u^c)$ while $\Pi^r(w, p)$ approaches 0. ∎

Observe that in Example 1, the heuristic wholesale price, $w^{1/2}$, for the supplier corresponds to optimal Stackelberg wholesale price when $\delta < 0$. Thus, the supplier cannot guarantee more than $\frac{1}{2}$ of the global optimal profit by acting selfishly. In addition, the example also illustrates that when the supplier acts selfishly, the retailers profit may be exactly zero, and hence the total supply chain profit can be as low as half of the global optimization profit. An important corollary, therefore, of Theorems 3.23 and 3.28 is as follows:

Corollary 3.29. *The worst-case system efficiency for both the competing retailers and the coordinated retailers systems in a regime with retail price caps, when the supplier uses the suggested buy-back menu scheme, is 50% of the global optimal supply chain profit. Additionally, in the worst-case, it is possible that the retailers make no profit.*

The following theorem generalizes the results of Theorem 3.28 to show that under any proposed buy-back menu scheme, the supplier can gain a higher expected profit with competing retailers rather than coordinated ones, while the customers experience the same retail prices and service levels in either setting.

Theorem 3.30. *For any wholesale price vector $\mathbf{w}^o \leq \mathbf{p}^u$ and any buy-back menu $\mathbf{s}^o(\mathbf{w}^o, \bullet)$ applied to the coordinated system, there exists a corresponding wholesale price vector $\mathbf{w}^d$ and a buy-back menu $\mathbf{s}^d(\mathbf{w}^d, \bullet)$ in the competing retailers system such that $\Pi_d^s(\mathbf{w}^d, \mathbf{s}^d(\mathbf{w}^d, \bullet)) \geq \Pi_o^s(\mathbf{w}^o, \mathbf{s}^o(\mathbf{w}^o, \bullet))$. In addition, the retailer prices, order quantities and hence the expected system-wide profit are identical for both systems.*

141

See proof in Appendix A.0.9

3.7. Concluding Remarks

In this work, we analyze the benefit of buy-back menu contracts for supply chains with a single supplier and competing retailers. We show that a selfish supplier can lock in at least 50% of the global optimization supply chain profit. In fact, in this case, when the supplier behaves selfishly, the retailers may have no profit at all. We provide a remedy wholesale price strategy that still guarantees the supplier 50% of the global profit but increases supply chain profit to at least 75% of the global optimal supply chain profit. In this case, retailers are guaranteed to lock in some profit, as long as retail prices are not at their upper bounds.

In addition, our work provides certain normative insights into supply chain operations. One of our key results is that the supplier's expected profit in a purely decentralized supply chain is higher than in a supply chain with coordinated retailers. This is true even when retail prices and customer service level are identical in both systems. Similarly, for the same supplier strategy, customers will face a lower retail price in a decentralized supply chain.

A corollary of our results arises when, due to either governmental or market restrictions or due to impositions of the supplier, there are upper bounds on the retail prices that retailers might set. Our work indicates that in such cases, there may exist circumstances where it is optimal for the supplier to set its wholesale price at this upper bound, leaving the retailer with no profit whatsoever. In such cases, the phenomenon of double marginalization vanishes, and since there is no incentive for the retailer to remain in the system, this may lead to the collapse of a supply chain to a single entity. Moreover as discussed in the previous section, from a system perspective, an argument may be made against the setting of price caps, since the use of price caps increases the price of anarchy of supply chains.

Appendix A

Appendix to "The Impact of Pricing and Buy-back Menus on Supply Chain Performance"

A.0.1 Proof of Observation 3.2

Proof. Consider a single-supplier single-retailer system similar to the one described earlier, except with stochasticity added into the demand model, i.e., $d = a - bp + \varepsilon$. Suppose that $a = 2 + \delta, b = 1$, the production cost $c = 0$, and let $p^u = 2 + \delta$ be an upper bound on the price that the retailer can charge and where δ is a small positive number.

The density function of ε is given by: $P\left(4^k\right) = \frac{1}{4^{k+2}}$ for $k = 0, 1, 2, ..., 31$ and $P(0) = 1 - \sum_{k=0}^{31} p\left(4^k\right)$. It follows that $\mu = E(\varepsilon) = \sum_{k=0}^{31} 4^k P\left(4^k\right) = 32\left(\frac{1}{16}\right) = 2$.

Consider the centralized system. Clearly, the system needs to produce at least an amount to cover the "deterministic" portion of the demand, i.e., an amount equal to $a - bp$. In addition, it needs to add an amount u to cover for the uncertain component of customer demand.

Thus, the optimal global expected profit is easily seen to be $\Pi^c(p, u) = p(a - bp + \mu) = \Pi^c(p, 4^{31})$, since production cost is zero. Moreover, $\Pi^c(p, 4^{31})$ attains its maximum value of

$\frac{(a+\mu)^2}{4b} = \frac{(4+\delta)^2}{4}$ at $p = \frac{a+\mu}{2b}$.

In the decentralized system, the retailer needs to order an amount equal to $a - bp$ plus an amount u to cover for the uncertain demand. The retailer's expected profit is thus given by

$$\Pi^r(w, p, u) = p\left(a - bp + E\left[u - \varepsilon\right]^+\right) - w\left(a - bp + u\right).$$

When $u = 0$, $\Pi^r(w, p, 0) = (p - w)(a - bp)$ and it attains its maximum value of $\frac{(a-bw)^2}{4b}$ at $p = \frac{a+bw}{2b}$.

The supplier's profit in the decentralized system is given by $\Pi^s(w, p, u) = w(a - bp + u)$. When $u = 0$, $\Pi^s\left(w, \frac{a+bw}{2b}, 0\right) = \frac{1}{2}w(a - bw)$ with maximum value of $\frac{a^2}{8b} = \frac{(2+\delta)^2}{8}$ at $w = \frac{a}{2b}$. The retailer's expected profit is $\frac{(2+\delta)^2}{16}$.

If however $u > 0$, then it follows from the optimality of the u^{th} unit that $pP\{\varepsilon \geq u\} \geq w$. Since $P\{\varepsilon > 0\} = \sum_{k=0}^{31} P(4^k) = \frac{1}{16}\sum_{k=0}^{31}\frac{1}{4^k} < \frac{1}{16}\sum_{k=0}^{\infty}\frac{1}{4^k} = \frac{1}{12}$, $w \leq \frac{p^u}{12}$ and $w(a - bp) < wa < \frac{p^u}{12}a = \frac{(2+\delta)^2}{12}$. Also, $wu \leq pP\{\varepsilon \geq u\}u = pP(u)u\sum_{v \geq u}\frac{P(v)}{P(u)} < puP(u)\sum_{k=0}^{\infty}\frac{1}{4^k} = \frac{p}{12} \leq \frac{p^u}{12} = \frac{2+\delta}{12}$. Hence, for $u > 0$, $\Pi^s(w, p, u) = w(a - bp + u) < \frac{(2+\delta)^2}{12} + \frac{2+\delta}{12} = \frac{6+5\delta+\delta^2}{12} < \frac{4+4\delta+\delta^2}{8} = \Pi^s\left(\frac{a}{2b}, \frac{3a}{4b}, 0\right)$. Hence, the supplier attains its maximum profit at $u = 0$ and so correspondingly, by setting $w = \frac{a}{2b}$, induces $u = 0$ and $p = \frac{3a}{4b}$ as the optimal decision for the retailer.

The ratio of the optimal expected supplier profit and the optimal expected system profit would then be $\frac{(2+\delta)^2}{8} / \frac{(4+\delta)^2}{4} \to \frac{1}{8}$ as $\delta \to 0$. Similarly, the ratio of the corresponding system expected profit and the optimal system expected profit $= \frac{3(2+\delta)^2}{16} / \frac{(4+\delta)^2}{4} \to \frac{3}{16}$ as $\delta \to 0$. $\blacksquare$

A.0.2 Proof of Lemma 3.9

Proof. The proof is by induction. We first show that $\mathbf{p}_{\overline{L_0}}^0 < \mathbf{p}_{\overline{L_0}}^u$. Observe that by rewriting $\tilde{\mathbf{a}}_{\overline{L_0}}$ as:

$$\tilde{\mathbf{a}}_{\overline{L_0}} = \left(\tilde{\mathbf{B}}\tilde{\mathbf{B}}^{-1}\tilde{\mathbf{a}}\right)_{\overline{L_0}} = \tilde{\mathbf{B}}_{\overline{L_0}}\left(\tilde{\mathbf{B}}^{-1}\tilde{\mathbf{a}}\right)_{\overline{L_0}} + \tilde{\mathbf{B}}_{\overline{L_0},L_0}\left(\tilde{\mathbf{B}}^{-1}\tilde{\mathbf{a}}\right)_{L_0} ,$$

we have that:

$$\widetilde{\mathbf{B}}_{\overline{L_0}}^{-1}\widetilde{\mathbf{a}}_{\overline{L_0}} = \left(\widetilde{\mathbf{B}}^{-1}\widetilde{\mathbf{a}}\right)_{\overline{L_0}} + \widetilde{\mathbf{B}}_{\overline{L_0}}^{-1}\widetilde{\mathbf{B}}_{\overline{L_0},L_0}\left(\widetilde{\mathbf{B}}^{-1}\widetilde{\mathbf{a}}\right)_{L_0} \tag{A.1}$$

$$
\begin{aligned}
\text{Now,} \quad \mathbf{p}_{\overline{L_0}}^0 &= \widetilde{\mathbf{B}}_{\overline{L_0}}^{-1}\left(\widetilde{\mathbf{a}}_{\overline{L_0}} - \widetilde{\mathbf{B}}_{\overline{L_0},L_0}\mathbf{p}_{L_0}^u\right) \\
&= \widetilde{\mathbf{B}}_{\overline{L_0}}^{-1}\widetilde{\mathbf{a}}_{\overline{L_0}} - \widetilde{\mathbf{B}}_{\overline{L_0}}^{-1}\widetilde{\mathbf{B}}_{\overline{L_0},L_0}\left(\widetilde{\mathbf{B}}^{-1}\widetilde{\mathbf{a}}\right)_{L_0} + \widetilde{\mathbf{B}}_{\overline{L_0}}^{-1}\widetilde{\mathbf{B}}_{\overline{L_0},L_0}\left[\left(\widetilde{\mathbf{B}}^{-1}\widetilde{\mathbf{a}}\right)_{L_0} - \mathbf{p}_{L_0}^u\right] \\
&\leq \widetilde{\mathbf{B}}_{\overline{L_0}}^{-1}\widetilde{\mathbf{a}}_{\overline{L_0}} - \widetilde{\mathbf{B}}_{\overline{L_0}}^{-1}\widetilde{\mathbf{B}}_{\overline{L_0},L_0}\left(\widetilde{\mathbf{B}}^{-1}\widetilde{\mathbf{a}}\right)_{L_0} \tag{A.2} \\
&= \left(\widetilde{\mathbf{B}}^{-1}\widetilde{\mathbf{a}}\right)_{\overline{L_0}} < \mathbf{p}_{\overline{L_0}}^u. \tag{A.3}
\end{aligned}
$$

where inequality (A.2) follows by noting that (a) by **Step 1** of Algorithm A, $\left(\widetilde{\mathbf{B}}^{-1}\widetilde{\mathbf{a}}\right)_{L_0} \geq \mathbf{p}_{L_0}^u$, (b) by assumption 1, all elements of $\widetilde{\mathbf{B}}_{\overline{L_0},L_0}$ are non-positive, and (c) according to Observation 3.4, all entries of $\widetilde{\mathbf{B}}_{\overline{L_0}}^{-1}$ are non-negative. Finally, inequality (A.3) follows from equation (A.1) and **Step 1** of Algorithm A.

Applying the inductive hypothesis, now assume that $\mathbf{p}_{\overline{L_k}}^k < \mathbf{p}_{\overline{L_k}}^u$. From **Step 2** of the Algorithm, we have that for each $i \in L_k \backslash L_{k+1} = \overline{L_{k+1}}\backslash\overline{L_k}$,

$$\frac{1}{\widetilde{b}_{ii}}\left[\widetilde{a}_i - \widetilde{\mathbf{B}}_{\{i\},L_k\backslash\{i\}}\mathbf{p}_{L_k\backslash\{i\}}^u - \widetilde{\mathbf{B}}_{\{i\},\overline{L_k}}\mathbf{p}_{\overline{L_k}}^k\right] < \mathbf{p}_i^u$$

Rearranging and combining the above inequality for all $i \in \overline{L_{k+1}}\backslash\overline{L_k}$, we have that:

$$\widetilde{\mathbf{a}}_{\overline{L_{k+1}}\backslash\overline{L_k}} < \widetilde{\mathbf{B}}_{\overline{L_{k+1}}\backslash\overline{L_k},L_k}\mathbf{p}_{L_k}^u + \widetilde{\mathbf{B}}_{\overline{L_{k+1}}\backslash\overline{L_k},\overline{L_k}}\mathbf{p}_{\overline{L_k}}^k = \widetilde{\mathbf{B}}_{\overline{L_{k+1}}\backslash\overline{L_k},N}\mathbf{p}^k \tag{A.4}$$

Similarly from the definition of $\mathbf{p}_{\overline{L_k}}^k$ we have that:

$$\widetilde{\mathbf{a}}_{\overline{L_k}} = \widetilde{\mathbf{B}}_{\overline{L_k}}\mathbf{p}_{\overline{L_k}}^k + \widetilde{\mathbf{B}}_{\overline{L_k},L_k}\mathbf{p}_{L_k}^u = \widetilde{\mathbf{B}}_{\overline{L_k}}\mathbf{p}_{\overline{L_k}}^k + \widetilde{\mathbf{B}}_{\overline{L_k},L_k}\mathbf{p}_{L_k}^k = \widetilde{\mathbf{B}}_{\overline{L_k},N}\mathbf{p}^k \tag{A.5}$$

145

Combining equations (A.4) and (A.5), and observing that $L_{k+1} \subseteq L_k$ we have that:

$$
\begin{aligned}
\mathbf{p}_{\overline{L_{k+1}}}^{k+1} &= \widetilde{\mathbf{B}}_{\overline{L_{k+1}}}^{-1} \left(\widetilde{\mathbf{a}}_{\overline{L_{k+1}}} - \widetilde{\mathbf{B}}_{\overline{L_{k+1}},L_{k+1}} \mathbf{p}_{L_{k+1}}^{u} \right) \\
&\leq \widetilde{\mathbf{B}}_{\overline{L_{k+1}}}^{-1} \left(\widetilde{\mathbf{B}}_{\overline{L_{k+1}},N} \mathbf{p}^{k} - \widetilde{\mathbf{B}}_{\overline{L_{k+1}},L_{k+1}} \mathbf{p}_{L_{k+1}}^{k} \right) \qquad (A.6) \\
&= \widetilde{\mathbf{B}}_{\overline{L_{k+1}}}^{-1} \widetilde{\mathbf{B}}_{\overline{L_{k+1}}} \mathbf{p}_{\overline{L_{k+1}}}^{k} = \mathbf{p}_{\overline{L_{k+1}}}^{k} \leq \mathbf{p}_{\overline{L_{k+1}}}^{u}
\end{aligned}
$$

Notice that in the above, we also use the fact that all entries of $\widetilde{\mathbf{B}}_{\overline{L_{k+1}}}^{-1}$ are non-negative. Moreover, Observation 3.4 also implies that the diagonal entries of $\widetilde{\mathbf{B}}_{\overline{L_{k+1}}}^{-1}$ are positive. Combining this observation with the strict inequality of (A.4), we have further from the above inequality that $\mathbf{p}_{\overline{L_{k+1}} \setminus L_k}^{k+1} < \mathbf{p}_{\overline{L_{k+1}} \setminus L_k}^{k}$. Hence by induction, $\mathbf{p}_{\overline{L_k}}^{k} < \mathbf{p}_{\overline{L_k}}^{u}$. ∎

A.0.3 Proof of Lemma 3.13

Proof. That $\Pi^r(\mathbf{c}) \geq \Pi^r(\mathbf{c}, \mathbf{p}^K)$ follows from the fact that $\Pi^r(\mathbf{c})$ is the optimal value over a larger feasible region. By an analogous argument, $\Pi^r(\mathbf{c}, \mathbf{p}^K) \geq \Pi^r(\mathbf{c}, \mathbf{p}^c)$. Since $s_i^H(c_i, p_i^c) o_i^c = \left[p_i^c - \frac{p_i^c - c_i}{F_i(u_i^c)} \right] o_i^c = [p_i^c - (p_i^c - e_i)] o_i^c = e_i o_i^c$, we now have that:

$$
\begin{aligned}
\Pi^r(\mathbf{c}, \mathbf{p}^c) &= (\mathbf{p}^c - \mathbf{c})^t (\mathbf{a} - \mathbf{B}\mathbf{p}^c + \mathbf{u}^c) - \left[\mathbf{p}^c - S^H(\mathbf{c}, \mathbf{p}^c) \right]^t \mathbf{o}^c \\
&= (\mathbf{p}^c - \mathbf{c})^t (\mathbf{a} - \mathbf{B}\mathbf{p}^c + \mathbf{u}^c) - [\mathbf{p}^c - \mathbf{e}]^t \mathbf{o}^c \\
&= \Pi^c(\mathbf{p}^c, \mathbf{u}^c).
\end{aligned}
$$

It follows that $\Pi^r(\mathbf{c}) \geq \Pi^r(\mathbf{c}, \mathbf{p}^K) \geq \Pi^c(\mathbf{p}^c, \mathbf{u}^c)$.

To show that $\mathbf{p}^K \leq \mathbf{p}^c$, we provide a proof by contradiction. Suppose that $\mathbf{p}^K \not\leq \mathbf{p}^c$, and let P be the largest subset of N with $\mathbf{p}_P^K > \mathbf{p}_P^c$. Note that $P \subseteq \overline{K}$. Since $\mathbf{p}^K$ is the optimal solution to $\Pi^r(\mathbf{c}, \mathbf{p}^K)$, it follows that $\mathbf{p}_P^K$ is the optimal solution to maximizing $\Pi^r(\mathbf{c}, \mathbf{p})$ subject to the constraints $\mathbf{p}_{\overline{P}} = \mathbf{p}_{\overline{P}}^K$. Now, the expression for $\Pi^r(\mathbf{c}, \mathbf{p})$ subject to the constraints $\mathbf{p}_{\overline{P}} = \mathbf{p}_{\overline{P}}^K$ may be written as:

$$
\begin{aligned}
&(\mathbf{p}_P - \mathbf{c}_P)^t \left[\mathbf{a}_P - \mathbf{B}_P \mathbf{p}_P - \mathbf{B}_{P,\overline{P}} \mathbf{p}_{\overline{P}}^K + \mathbf{u}_P^c \right] - \left[\mathbf{p}_P - s_P^H(\mathbf{c}_P, \mathbf{p}_P) \right]^t \mathbf{o}_P^c \\
&+ \left[\mathbf{p}_{\overline{P}}^K - \mathbf{c}_{\overline{P}} \right]^t \left[\mathbf{a}_{\overline{P}} - \mathbf{B}_{\overline{P}} \mathbf{p}_{\overline{P}}^K - \mathbf{B}_{\overline{P},P} \mathbf{p}_P + \mathbf{u}_{\overline{P}}^c \right] - \left[\mathbf{p}_{\overline{P}}^K - s_{\overline{P}}^H(\mathbf{c}_{\overline{P}}, \mathbf{p}_{\overline{P}}^K) \right]^t \mathbf{o}_{\overline{P}}^c
\end{aligned}
$$

146

which is strictly concave in $\mathbf{p}_P$. As a consequence, it's derivative with respect to $\mathbf{p}_P$ at $\mathbf{p}_P = \mathbf{p}_P^c < \mathbf{p}_P^K$ must be positive. That is:

$$\mathbf{a}_P - \mathbf{B}_P \mathbf{p}_P^c - \mathbf{B}_{P,\overline{P}} \mathbf{p}_{\overline{P}}^K + \mathbf{u}_P^c - \mathbf{B}_P \left(\mathbf{p}_P^c - \mathbf{c}_P \right)$$

$$-\mathbf{B}_{P,\overline{P}} \left(\mathbf{p}_{\overline{P}}^K - \mathbf{c}_{\overline{P}} \right) - \mathbf{o}_P^c + \mathbf{s}_P^{H\prime} \left(\mathbf{c}_P, \mathbf{p}_P^c \right)^t \mathbf{o}_P^c \; > \; 0$$

Since $\mathbf{p}_{\overline{P}}^K \leq \mathbf{p}_{\overline{P}}^c$ and $\mathbf{s}_P^H \left(\mathbf{c}_P, \bullet \right)$ is a decreasing function of $\mathbf{p}_P$, we have that:

$$\mathbf{a}_P - \mathbf{B}_P \mathbf{p}_P^c - \mathbf{B}_{P,\overline{P}} \mathbf{p}_{\overline{P}}^c + \mathbf{u}_P^c$$

$$-\mathbf{B}_P \left(\mathbf{p}_P^c - \mathbf{c}_P \right) - \mathbf{B}_{P,\overline{P}} \left(\mathbf{p}_{\overline{P}}^c - \mathbf{c}_{\overline{P}} \right) - \mathbf{o}_P^c$$

$$\geq \; \mathbf{a}_P - \mathbf{B}_P \mathbf{p}_P^c - \mathbf{B}_{P,\overline{P}} \mathbf{p}_{\overline{P}}^K + \mathbf{u}_P^c$$

$$-\mathbf{B}_P \left(\mathbf{p}_P^c - \mathbf{c}_P \right) - \mathbf{B}_{P,\overline{P}} \left(\mathbf{p}_{\overline{P}}^K - \mathbf{c}_{\overline{P}} \right) - \mathbf{o}_P^c + \mathbf{s}_P^{H\prime} \left(\mathbf{c}_P, \mathbf{p}_P^c \right)^t \mathbf{o}_P^c$$

$$> \; 0. \tag{A.7}$$

On the other hand, we have that $\mathbf{p}^c$ maximizes the centralized system profit, $\Pi^c(\mathbf{p}, \mathbf{u}^c) = [\mathbf{p} - \mathbf{c}]^t [\mathbf{a} - \mathbf{B}\mathbf{p} + \mathbf{u}^c] - [\mathbf{p} - \mathbf{e}]^t \mathbf{o}^c$. Consider now $\Pi^c(\mathbf{p}, \mathbf{u}^c)$ subject to the constraint that $\mathbf{p}_{\overline{P}} = \mathbf{p}_{\overline{P}}^c$, which is the expression:

$$[\mathbf{p}_P - \mathbf{c}_P]^t \left[\mathbf{a}_P - \mathbf{B}_P \mathbf{p}_P - \mathbf{B}_{P,\overline{P}} \mathbf{p}_{\overline{P}}^c + \mathbf{u}_P^c \right] - [\mathbf{p}_P^c - \mathbf{e}_P]^t \mathbf{o}_P^c$$

$$+ \left[\mathbf{p}_{\overline{P}}^c - \mathbf{c}_{\overline{P}} \right]^t \left[\mathbf{a}_{\overline{P}} - \mathbf{B}_{\overline{P}} \mathbf{p}_{\overline{P}}^c - \mathbf{B}_{\overline{P},P} \mathbf{p}_P + \mathbf{u}_{\overline{P}}^c \right] - \left[\mathbf{p}_{\overline{P}}^c - \mathbf{e}_{\overline{P}} \right]^t \mathbf{o}_{\overline{P}}^c$$

Now, since the above expression is strictly concave in $\mathbf{p}_P$ and maximized at $\mathbf{p}_P^c (< \mathbf{p}_P^K \leq \mathbf{p}_P^u)$, it follows that it's derivative with respect to $\mathbf{p}_K$ at $\mathbf{p}_K = \mathbf{p}_K^c$ is zero. That is,

$$\mathbf{a}_P - \mathbf{B}_P \mathbf{p}_P^c - \mathbf{B}_{P,\overline{P}} \mathbf{p}_{\overline{P}}^c + \mathbf{u}_P^c - \mathbf{B}_P \left(\mathbf{p}_P^c - \mathbf{c}_P \right) - \mathbf{B}_{P,\overline{P}} \left(\mathbf{p}_{\overline{P}}^c - \mathbf{c}_{\overline{P}} \right) - \mathbf{o}_P^c = 0,$$

which contradicts equation (A.7). Hence, $\mathbf{p}^K \leq \mathbf{p}^c$. In particular, $\mathbf{p}(\mathbf{c}) = \mathbf{p}^\varnothing \leq \mathbf{p}^c$.

It remains to show that $\mathbf{p}(\mathbf{c}) \leq \mathbf{p}^K$. Observe that the derivative of $\Pi^r(\mathbf{c}, \mathbf{p}) = (\mathbf{p} - \mathbf{c})^t (\mathbf{a} - \mathbf{B}\mathbf{p} + \mathbf{u}^c - \mathbf{v}^c)$ with respect to $\mathbf{p}_{\overline{K}}$, given by:

$$\mathbf{a}_{\overline{K}} + \mathbf{u}_{\overline{K}}^c - \mathbf{v}_{\overline{K}}^c - \mathbf{B}_{\overline{K}} \mathbf{p}_{\overline{K}} - \mathbf{B}_{\overline{K},K} \mathbf{p}_K - \mathbf{B}_{\overline{K}} \left(\mathbf{p}_{\overline{K}} - \mathbf{c}_{\overline{K}} \right) - \mathbf{B}_{\overline{K},K} \left(\mathbf{p}_K - \mathbf{c}_K \right)$$

$$= \mathbf{a}_{\overline{K}} + \mathbf{u}_{\overline{K}}^c - \mathbf{v}_{\overline{K}}^c - 2\mathbf{B}_{\overline{K}} \mathbf{p}_{\overline{K}} - 2\mathbf{B}_{\overline{K},K} \mathbf{p}_K + \mathbf{B}_{\overline{K},N} \mathbf{c}$$

is a non-decreasing function of $\mathbf{p}_K$. Consider the optimal solution $\mathbf{p}^K$ of $\Pi^r \left(\mathbf{c}, \mathbf{p}^K \right)$, and

147

suppose that $L^K = \{i \in N : p_i^K = p_i^u\}$. Then by the KKT optimality conditions, it follows that:

$$
\begin{aligned}
\mathbf{p}^u_{\overline{K}\backslash L^K} \;>\; & \mathbf{p}^K_{\overline{K}\backslash L^K} \\
=\;& \tfrac{1}{2}\mathbf{B}^{-1}_{\overline{K}\backslash L^K}\left(\mathbf{a}_{\overline{K}\backslash L^K} + \mathbf{u}^c_{\overline{K}\backslash L^K} - \mathbf{v}^c_{\overline{K}\backslash L^K} + \mathbf{B}_{\overline{K}\backslash L^K,N}\mathbf{c}\right) \\
& -\mathbf{B}^{-1}_{\overline{K}\backslash L^K}\mathbf{B}_{\overline{K}\backslash L^K,L^K\backslash K}\mathbf{p}^u_{L^K\backslash K} - \mathbf{B}^{-1}_{\overline{K}\backslash L^K}\mathbf{B}_{\overline{K}\backslash L^K,K}\mathbf{p}^c_K \\
\geq\;& \tfrac{1}{2}\mathbf{B}^{-1}_{\overline{K}\backslash L^K}\left(\mathbf{a}_{\overline{K}\backslash L^K} + \mathbf{u}^c_{\overline{K}\backslash L^K} - \mathbf{v}^c_{\overline{K}\backslash L^K} + \mathbf{B}_{\overline{K}\backslash L^K,N}\mathbf{c}\right) \\
& -\mathbf{B}^{-1}_{\overline{K}\backslash L^K}\mathbf{B}_{\overline{K}\backslash L^K,L^K\backslash K}\mathbf{p}^u_{L^K\backslash K} - \mathbf{B}^{-1}_{\overline{K}\backslash L^K}\mathbf{B}_{\overline{K}\backslash L^K,K}\mathbf{p}_K\left(\mathbf{c}\right) \qquad\text{(A.8)} \\
\geq\;& \tfrac{1}{2}\mathbf{B}^{-1}_{\overline{K}\backslash L^K}\left(\mathbf{a}_{\overline{K}\backslash L^K} + \mathbf{u}^c_{\overline{K}\backslash L^K} - \mathbf{v}^c_{\overline{K}\backslash L^K} + \mathbf{B}_{\overline{K}\backslash L^K,N}\mathbf{c}\right) \\
& -\mathbf{B}^{-1}_{\overline{K}\backslash L^K}\mathbf{B}_{\overline{K}\backslash L^K,L^K\backslash K}\mathbf{p}_{L^K\backslash K}(\mathbf{c}) - \mathbf{B}^{-1}_{\overline{K}\backslash L^K}\mathbf{B}_{\overline{K}\backslash L^K,K}\mathbf{p}_K\left(\mathbf{c}\right) \qquad\text{(A.9)}
\end{aligned}
$$

where inequality (A.8) is implied by $\mathbf{p}\left(\mathbf{c}\right) \leq \mathbf{p}^c$ and inequality (A.9) is implied by $\mathbf{p}_{L^K}(\mathbf{c}) \leq \mathbf{p}^u_{L^K}$. Hence $\mathbf{p}^u_{\overline{K}\backslash L^K} > \mathbf{p}^K_{\overline{K}\backslash L^K} \geq \mathbf{p}_{\overline{K}\backslash L^K}\left(\mathbf{c}\right)$. Adding this to the facts that $\mathbf{p}^K_K = \mathbf{p}^c_K \geq \mathbf{p}_K\left(\mathbf{c}\right)$, and $\mathbf{p}^K_{L^K} = \mathbf{p}^u_{L^K} \geq \mathbf{p}_{L^K}\left(\mathbf{c}\right)$, we have the desired result that $\mathbf{p}\left(\mathbf{c}\right) \leq \mathbf{p}^K \leq \mathbf{p}^c$. ∎

A.0.4   Proof of Lemma 3.18

Proof. We prove the lemma via induction on i. For any $\rho_1 \leq \rho < \rho_2$, we first need to show that the desired bound holds for $\left(\mathbf{p}^\rho - \mathbf{c}\right)\left[\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c - \mathbf{B}\mathbf{p}^\rho\right]^t$, which may be rewritten as:

$$
\left(\mathbf{p}^\rho_{L(\rho)} - \mathbf{c}_{L(\rho)}\right)^t \left(\mathbf{a}_{L(\rho)} + \mathbf{u}^c_{L(\rho)} - \mathbf{v}^c_{L(\rho)} - \mathbf{B}_{L(\rho),N}\mathbf{p}^\rho\right) \qquad\text{(A.10)}
$$

$$
+ \left[\mathbf{p}^\rho_{\overline{L(\rho_1)}} - \mathbf{c}_{\overline{L(\rho_1)}}\right]^t \left[\mathbf{a}_{\overline{L(\rho_1)}} + \mathbf{u}^c_{\overline{L(\rho_1)}} - \mathbf{v}^c_{\overline{L(\rho_1)}} - \mathbf{B}_{\overline{L(\rho_1)},N}\mathbf{p}^\rho\right] \qquad\text{(A.11)}
$$

To provide these bounds, we will need to eliminate $\mathbf{p}^\rho$ from the above expression and instead replace it by $\mathbf{p}(\mathbf{c})$. Before we begin bounding these replacements in expressions (A.10) and (A.11), we present a few facts useful in this analysis.

148

Rearranging the optimality condition presented in equation (3.9), it may be shown that:

$$
\begin{aligned}
\mathbf{p}^{\rho}_{\overline{L(\rho)}} - \mathbf{p}_{\overline{L(\rho)}}\left(\mathbf{c}\right) &= \rho \mathbf{B}^{-1}_{\overline{L(\rho_1)}}\left[\mathbf{a}_{\overline{L(\rho_1)}} + \mathbf{u}^{c}_{\overline{L(\rho_1)}} - \mathbf{v}^{c}_{\overline{L(\rho_1)}} - \mathbf{B}_{\overline{L(\rho_1)},N}\mathbf{p}\left(\mathbf{c}\right)\right] \\
&\quad -\mathbf{B}^{-1}_{\overline{L(\rho_1)}}\mathbf{B}_{\overline{L(\rho_1)},L(\rho_1)}\left[\mathbf{p}^{\rho}_{L(\rho_1)} - \mathbf{p}_{L(\rho_1)}\left(\mathbf{c}\right)\right] \\
&= \rho \mathbf{B}^{-1}_{\overline{L(\rho_1)}}\left[\mathbf{a}_{\overline{L(\rho_1)}} + \mathbf{u}^{c}_{\overline{L(\rho_1)}} - \mathbf{v}^{c}_{\overline{L(\rho_1)}} - \mathbf{B}_{\overline{L(\rho_1)},N}\mathbf{p}\left(\mathbf{c}\right)\right] \quad\quad (A.12) \\
&= \rho \mathbf{B}^{-1}_{\overline{L(\rho_1)}}\mathbf{B}_{\overline{L(\rho_1)},N}\left[\mathbf{p}\left(\mathbf{c}\right) - \mathbf{c}\right] \quad\quad (A.13) \\
&= \rho\left[\mathbf{p}_{\overline{L(\rho_1)}}\left(\mathbf{c}\right) - \mathbf{c}_{\overline{L(\rho_1)}}\right] \\
&\quad +\rho\mathbf{B}^{-1}_{\overline{L(\rho_1)}}\mathbf{B}_{\overline{L(\rho_1)},L(\rho_1)}\left[\mathbf{p}_{L(\rho_1)}\left(\mathbf{c}\right) - \mathbf{c}_{L(\rho_1)}\right] \quad\quad (A.14)
\end{aligned}
$$

Equation (A.12) follows from the fact that $L(\rho_1) = L(\rho)$ and $\mathbf{p}^{\rho_1} = \mathbf{p}(\mathbf{c})$, which in turn imply that $\mathbf{p}^{\rho}_{L(\rho)} = \mathbf{p}_{L(\rho)}(\mathbf{c})$. Equation (A.13) is derived using the KKT optimality conditions for $\mathbf{p}(\mathbf{c})$ in the indices $\overline{L(\rho_1)}$. Finally, equation (A.14) is obtained by expanding out $\mathbf{B}_{\overline{L(\rho_1)},N}$.

Consider the expression (A.10). We now have that:

$$
\begin{aligned}
&\left(\mathbf{p}^{\rho}_{L(\rho)} - \mathbf{c}_{L(\rho)}\right)^{t}\left(\mathbf{a}_{L(\rho)} + \mathbf{u}^{c}_{L(\rho)} - \mathbf{v}^{c}_{L(\rho)} - \mathbf{B}_{L(\rho),N}\mathbf{p}^{\rho}\right) \\
&= \left[\mathbf{p}_{L(\rho_1)}\left(\mathbf{c}\right) - \mathbf{c}_{L(\rho_1)}\right]^{t}\left(\mathbf{a}_{L(\rho_1)} + \mathbf{u}^{c}_{L(\rho_1)} - \mathbf{v}^{c}_{L(\rho_1)} - \mathbf{B}_{L(\rho_1),N}\mathbf{p}^{\rho}\right) \\
&= \left[\mathbf{p}_{L(\rho_1)}\left(\mathbf{c}\right) - \mathbf{c}_{L(\rho_1)}\right]^{t}\left[\mathbf{a}_{L(\rho_1)} + \mathbf{u}^{c}_{L(\rho_1)} - \mathbf{v}^{c}_{L(\rho_1)} - \mathbf{B}_{L(\rho_1),N}\mathbf{p}\left(\mathbf{c}\right)\right] \\
&\quad - \left[\mathbf{p}_{L(\rho_1)}\left(\mathbf{c}\right) - \mathbf{c}_{L(\rho_1)}\right]^{t}\mathbf{B}_{L(\rho_1),\overline{L(\rho_1)}}\left[\mathbf{p}^{\rho}_{\overline{L(\rho)}} - \mathbf{p}_{\overline{L(\rho)}}\left(\mathbf{c}\right)\right] \\
&= \left[\mathbf{p}_{L(\rho_1)}\left(\mathbf{c}\right) - \mathbf{c}_{L(\rho_1)}\right]^{t}\left[\mathbf{a}_{L(\rho_1)} + \mathbf{u}^{c}_{L(\rho_1)} - \mathbf{v}^{c}_{L(\rho_1)} - \mathbf{B}_{L(\rho_1),N}\mathbf{p}\left(\mathbf{c}\right)\right]
\end{aligned}
$$

$$
-\rho\left[\mathbf{p}_{L(\rho_1)}\left(\mathbf{c}\right) - \mathbf{c}_{L(\rho_1)}\right]^{t}\mathbf{B}_{L(\rho_1),\overline{L(\rho_1)}}\mathbf{B}^{-1}_{\overline{L(\rho_1)}}\left[\mathbf{a}_{\overline{L(\rho_1)}} + \mathbf{u}^{c}_{\overline{L(\rho_1)}} - \mathbf{v}^{c}_{\overline{L(\rho_1)}} - \mathbf{B}_{\overline{L(\rho_1)},N}\mathbf{p}\left(\mathbf{c}\right)\right]
$$
$$(A.15)$$

where equation (A.15) is obtained by substituting equation (A.12) in the above expression. Similarly from expression (A.11), we have that:

$$
\left[\mathbf{p}^{\rho}_{\overline{L(\rho_1)}} - \mathbf{c}_{\overline{L(\rho_1)}}\right]^{t}\left[\mathbf{a}_{\overline{L(\rho_1)}} + \mathbf{u}^{c}_{\overline{L(\rho_1)}} - \mathbf{v}^{c}_{\overline{L(\rho_1)}} - \mathbf{B}_{\overline{L(\rho_1)},N}\mathbf{p}^{\rho}\right]
$$

which may be rearranged as

$$
\begin{aligned}
&= \left[\mathbf{p}^{\rho}_{\overline{L(\rho_1)}} - \mathbf{p}_{\overline{L(\rho_1)}}(\mathbf{c})\right]^{t}\left[\mathbf{a}_{\overline{L(\rho_1)}} + \mathbf{u}^{c}_{\overline{L(\rho_1)}} - \mathbf{v}^{c}_{\overline{L(\rho_1)}} - \mathbf{B}_{\overline{L(\rho_1)},N}\mathbf{p}^{\rho}\right] \\
&\quad + \left[\mathbf{p}_{\overline{L(\rho_1)}}(\mathbf{c}) - \mathbf{c}_{\overline{L(\rho_1)}}\right]^{t}\left[\mathbf{a}_{\overline{L(\rho_1)}} + \mathbf{u}^{c}_{\overline{L(\rho_1)}} - \mathbf{v}^{c}_{\overline{L(\rho_1)}} - \mathbf{B}_{\overline{L(\rho_1)},N}\mathbf{p}^{\rho}\right]
\end{aligned}
$$

Substituting from the optimality condition given in equation (3.9), this may be written as:

149

$$= (1 - \rho) \left[\mathbf{p}^\rho_{\overline{L(\rho_1)}} - \mathbf{p}_{\overline{L(\rho_1)}}(\mathbf{c}) \right]^t \left[\mathbf{a}_{\overline{L(\rho_1)}} + \mathbf{u}^c_{\overline{L(\rho_1)}} - \mathbf{v}^c_{\overline{L(\rho_1)}} - \mathbf{B}_{\overline{L(\rho_1)},N}\mathbf{p}(\mathbf{c}) \right]$$

$$+ (1 - \rho) \left[\mathbf{p}_{\overline{L(\rho_1)}}(\mathbf{c}) - \mathbf{c}_{\overline{L(\rho_1)}} \right]^t \left[\mathbf{a}_{\overline{L(\rho_1)}} + \mathbf{u}^c_{\overline{L(\rho_1)}} - \mathbf{v}^c_{\overline{L(\rho_1)}} - \mathbf{B}_{\overline{L(\rho_1)},N}\mathbf{p}(\mathbf{c}) \right]$$

Substituting for $\left[\mathbf{p}^\rho_{\overline{L(\rho)}} - \mathbf{p}_{\overline{L(\rho)}}(\mathbf{c}) \right]$ from equation (A.14) in the above expression, we have that:

$$= (1 + \rho)(1 - \rho) \left[\mathbf{p}_{\overline{L(\rho_1)}}(\mathbf{c}) - \mathbf{c}_{\overline{L(\rho_1)}} \right]^t \left[\mathbf{a}_{\overline{L(\rho_1)}} + \mathbf{u}^c_{\overline{L(\rho_1)}} - \mathbf{v}^c_{\overline{L(\rho_1)}} - \mathbf{B}_{\overline{L(\rho_1)},N}\mathbf{p}(\mathbf{c}) \right]$$

$$+ \rho \left[\mathbf{p}_{L(\rho_1)}(\mathbf{c}) - \mathbf{c}_{L(\rho_1)} \right]^t \mathbf{B}_{L(\rho_1),\overline{L(\rho_1)}}\mathbf{B}^{-1}_{\overline{L(\rho_1)}} \left[\mathbf{a}_{\overline{L(\rho_1)}} + \mathbf{u}^c_{\overline{L(\rho_1)}} - \mathbf{v}^c_{\overline{L(\rho_1)}} - \mathbf{B}_{\overline{L(\rho_1)},N}\mathbf{p}(\mathbf{c}) \right]$$

$$- \rho^2 \left[\mathbf{p}_{L(\rho_1)}(\mathbf{c}) - \mathbf{c}_{L(\rho_1)} \right]^t \mathbf{B}_{L(\rho_1),\overline{L(\rho_1)}}\mathbf{B}^{-1}_{\overline{L(\rho_1)}} \left[\mathbf{a}_{\overline{L(\rho_1)}} + \mathbf{u}^c_{\overline{L(\rho_1)}} - \mathbf{v}^c_{\overline{L(\rho_1)}} - \mathbf{B}_{\overline{L(\rho_1)},N}\mathbf{p}(\mathbf{c}) \right]$$

Adding the above expression with equation (A.15), we have that:

$$(\mathbf{p}^\rho - \mathbf{c})[\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c - \mathbf{B}\mathbf{p}^\rho]^t$$

$$= \left[\mathbf{p}_{L(\rho_1)}(\mathbf{c}) - \mathbf{c}_{L(\rho_1)} \right]^t \left[\mathbf{a}_{L(\rho_1)} + \mathbf{u}^c_{L(\rho_1)} - \mathbf{v}^c_{L(\rho_1)} - \mathbf{B}_{L(\rho_1),N}\mathbf{p}(\mathbf{c}) \right]$$

$$+ (1 + \rho)(1 - \rho) \left[\mathbf{p}_{\overline{L(\rho_1)}}(\mathbf{c}) - \mathbf{c}_{\overline{L(\rho_1)}} \right]^t \left[\mathbf{a}_{\overline{L(\rho_1)}} + \mathbf{u}^c_{\overline{L(\rho_1)}} - \mathbf{v}^c_{\overline{L(\rho_1)}} - \mathbf{B}_{\overline{L(\rho_1)},N}\mathbf{p}(\mathbf{c}) \right]$$

$$- \rho^2 \left[\mathbf{p}_{L(\rho_1)}(\mathbf{c}) - \mathbf{c}_{L(\rho_1)} \right]^t \mathbf{B}_{L(\rho_1),\overline{L(\rho_1)}}\mathbf{B}^{-1}_{\overline{L(\rho_1)}} \left[\mathbf{a}_{\overline{L(\rho_1)}} + \mathbf{u}^c_{\overline{L(\rho_1)}} - \mathbf{v}^c_{\overline{L(\rho_1)}} - \mathbf{B}_{\overline{L(\rho_1)},N}\mathbf{p}(\mathbf{c}) \right]$$

$$\geq \left[\mathbf{p}_{L(\rho_1)}(\mathbf{c}) - \mathbf{c}_{L(\rho_1)} \right]^t \left[\mathbf{a}_{L(\rho_1)} + \mathbf{u}^c_{L(\rho_1)} - \mathbf{v}^c_{L(\rho_1)} - \mathbf{B}_{L(\rho_1),N}\mathbf{p}(\mathbf{c}) \right]$$

$$+ (1 + \rho)(1 - \rho) \left[\mathbf{p}_{\overline{L(\rho_1)}}(\mathbf{c}) - \mathbf{c}_{\overline{L(\rho_1)}} \right]^t \left[\mathbf{a}_{\overline{L(\rho_1)}} + \mathbf{u}^c_{\overline{L(\rho_1)}} - \mathbf{v}^c_{\overline{L(\rho_1)}} - \mathbf{B}_{\overline{L(\rho_1)},N}\mathbf{p}(\mathbf{c}) \right].$$

So, the statement holds for $i = 1$.

For any $\rho_{i-1} \leq \rho < \rho_i$ for $i = 2, 3, ..., \theta - 1$, assume $(\mathbf{p}^\rho - \mathbf{c})[\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c - \mathbf{B}\mathbf{p}^\rho]$

$$\geq \sum_{j=1}^{i-1} (1 + \rho_j)(1 - \rho_j) \left[\mathbf{p}_{\tilde{L}(\rho_j)}(\mathbf{c}) - \mathbf{c}_{\tilde{L}(\rho_j)} \right]^t \left[\mathbf{a}_{\tilde{L}(\rho_j)} + \mathbf{u}^c_{\tilde{L}(\rho_j)} - \mathbf{v}^c_{\tilde{L}(\rho_j)} - \mathbf{B}_{\tilde{L}(\rho_j),N}\mathbf{p}(\mathbf{c}) \right]$$

$$+ (1 + \rho)(1 - \rho) \left[\mathbf{p}_{\overline{L(\rho_{i-1})}}(\mathbf{c}) - \mathbf{c}_{\overline{L(\rho_{i-1})}} \right]^t \left[\mathbf{a}_{\overline{L(\rho_{i-1})}} + \mathbf{u}^c_{\overline{L(\rho_{i-1})}} - \mathbf{v}^c_{\overline{L(\rho_{i-1})}} - \mathbf{B}_{\overline{L(\rho_{i-1})},N}\mathbf{p}(\mathbf{c}) \right].$$

For any $\rho_i \leq \rho < \rho_{i+1}$ or $\rho_{\theta-1} \leq \rho \leq \rho_\theta$ when $i = \theta - 1$, we have to prove a bound on:

$$(\mathbf{p}^\rho - \mathbf{c})[\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c - \mathbf{B}\mathbf{p}^\rho] = (\mathbf{p}^{\rho_i} - \mathbf{c})[\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c - \mathbf{B}\mathbf{p}^{\rho_i}] \qquad (A.16)$$

$$+ (\mathbf{p}^\rho - \mathbf{c})[\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c - \mathbf{B}\mathbf{p}^\rho] - (\mathbf{p}^{\rho_i} - \mathbf{c})[\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c - \mathbf{B}\mathbf{p}^{\rho_i}] \qquad (A.17)$$

The expression in equation(A.17) can be further rewritten in two parts, corresponding to the sets $L(\rho_i)$ and $\overline{L(\rho_i)}$. We shall analyze each of these two parts separately. First, consider expression (A.17) for the set $L(\rho_i)$:

$$\left[\mathbf{p}^{\rho}_{L(\rho_i)} - \mathbf{c}_{L(\rho_i)}\right]^t \left[\mathbf{a}_{L(\rho_i)} + \mathbf{u}^c_{L(\rho_i)} - \mathbf{v}^c_{L(\rho_i)} - \mathbf{B}_{L(\rho_i),N}\mathbf{p}^{\rho}\right]$$
$$- \left[\mathbf{p}^{\rho_i}_{L(\rho_i)} - \mathbf{c}_{L(\rho_i)}\right]^t \left[\mathbf{a}_{L(\rho_i)} + \mathbf{u}^c_{L(\rho_i)} - \mathbf{v}^c_{L(\rho_i)} - \mathbf{B}_{L(\rho_i),N}\mathbf{p}^{\rho_i}\right]$$

From the equality of $L(\rho_i) = L(\rho)$, this may be simplified as:

$$= -\left[\mathbf{p}^{\rho_i}_{L(\rho_i)} - \mathbf{c}_{L(\rho_i)}\right]^t \mathbf{B}_{L(\rho_i),\overline{L(\rho_i)}} \left[\mathbf{p}^{\rho}_{\overline{L(\rho_i)}} - \mathbf{p}^{\rho_i^-}_{\overline{L(\rho_i)}}\right]$$

Substituting for the formulas of $\mathbf{p}^{\rho}_{\overline{L(\rho_i)}}$ and $\mathbf{p}^{\rho_i^-}_{\overline{L(\rho_i)}}$, the above may be rewritten as:

$$= -(\rho - \rho_i)\left[\mathbf{p}^{\rho_i}_{L(\rho_i)} - \mathbf{c}_{L(\rho_i)}\right]^t \mathbf{B}_{L(\rho_i),\overline{L(\rho_i)}}\mathbf{B}^{-1}_{\overline{L(\rho_i)}}\left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N}\mathbf{p}\,(\mathbf{c})\right]$$

(A.18)

Now, consider the expression (A.17) restricted to the set of indices $\overline{L(\rho_i)}$:

$$\left[\mathbf{p}^{\rho}_{\overline{L(\rho_i)}} - \mathbf{c}_{\overline{L(\rho_i)}}\right]^t \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N}\mathbf{p}^{\rho}\right]$$
$$- \left[\mathbf{p}^{\rho_i}_{\overline{L(\rho_i)}} - \mathbf{c}_{\overline{L(\rho_i)}}\right]^t \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N}\mathbf{p}^{\rho_i}\right]$$
$$= \left[\mathbf{p}^{\rho}_{\overline{L(\rho_i)}} - \mathbf{c}_{\overline{L(\rho_i)}}\right]^t \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N}\mathbf{p}^{\rho}\right]$$
$$- \left[\mathbf{p}^{\rho_i}_{\overline{L(\rho_i)}} - \mathbf{c}_{\overline{L(\rho_i)}}\right]^t \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N}\mathbf{p}^{\rho}\right]$$
$$+ \left[\mathbf{p}^{\rho_i}_{\overline{L(\rho_i)}} - \mathbf{c}_{\overline{L(\rho_i)}}\right]^t \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N}\mathbf{p}^{\rho}\right]$$
$$- \left[\mathbf{p}^{\rho_i}_{\overline{L(\rho_i)}} - \mathbf{c}_{\overline{L(\rho_i)}}\right]^t \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N}\mathbf{p}^{\rho_i}\right]$$

From the optimality conditions of equation (3.9) for ρ and ρ_i, the above expression reduces to:

$$= \left[\mathbf{p}^{\rho}_{\overline{L(\rho_i)}} - \mathbf{p}^{\rho_i}_{\overline{L(\rho_i)}}\right]^t \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N}\mathbf{p}^{\rho}\right]$$

(A.19)

$$- (\rho - \rho_i)\left[\mathbf{p}^{\rho_i}_{\overline{L(\rho_i)}} - \mathbf{c}_{\overline{L(\rho_i)}}\right]^t \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N}\mathbf{p}\,(\mathbf{c})\right]$$

(A.20)

In order the replace the $\mathbf{p}(\rho)$s by $\mathbf{p}(\mathbf{c})$ in expressions (A.19) and (A.20), we need to derive a few additional facts. From the optimality conditions of equation (3.9) for ρ and ρ_i, we have that:

$$\mathbf{p}^{\rho}_{\overline{L(\rho_i)}} - \mathbf{p}^{\rho_i}_{\overline{L(\rho_i)}}$$
$$= (\rho - \rho_i)\,\mathbf{B}^{-1}_{\overline{L(\rho_i)}}\left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N}\mathbf{p}\,(\mathbf{c})\right]$$
$$- \mathbf{B}^{-1}_{\overline{L(\rho_i)}}\mathbf{B}_{\overline{L(\rho_i)},L(\rho_i)}\left[\mathbf{p}^{\rho}_{L(\rho_i)} - \mathbf{p}^{\rho_i}_{L(\rho_i)}\right]$$

151

Since $L(\rho) = L(\rho_i)$ implies that $\mathbf{p}^{\rho}_{L(\rho_i)} = \mathbf{p}^{\rho_i}_{L(\rho_i)}$, and from the optimality of $\mathbf{p}(\mathbf{c})$ for wholesale price $\mathbf{c}$:

$$= (\rho - \rho_i)\, \mathbf{B}^{-1}_{\overline{L(\rho_i)}} \mathbf{B}_{\overline{L(\rho_i)},N} \left[\mathbf{p}(\mathbf{c}) - \mathbf{c}\right]$$

$$= (\rho - \rho_i)\left[\mathbf{p}_{\overline{L(\rho_i)}}(\mathbf{c}) - \mathbf{c}_{\overline{L(\rho_i)}}\right] + (\rho - \rho_i)\, \mathbf{B}^{-1}_{\overline{L(\rho_i)}} \mathbf{B}_{\overline{L(\rho_i)},L(\rho_i)} \left[\mathbf{p}_{L(\rho_i)}(\mathbf{c}) - \mathbf{c}_{L(\rho_i)}\right]. \tag{A.21}$$

Additionally, we have from the definition of $\mathbf{p}^{\rho_i}_{\overline{L(\rho_i)}}$ that:

$$\left(\mathbf{p}^{\rho_i}_{\overline{L(\rho_i)}} - \mathbf{p}_{\overline{L(\rho_i)}}(\mathbf{c})\right)$$

$$= \mathbf{B}^{-1}_{\overline{L(\rho_i)}}\left[\rho_i\left(\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}}\right) + (1-\rho_i)\mathbf{B}_{\overline{L(\rho_i)},N}\mathbf{p}(\mathbf{c}) - \mathbf{B}_{\overline{L(\rho_i)},L(\rho_i)}\mathbf{p}^u_{L(\rho_i)}\right] - \mathbf{p}_{\overline{L(\rho_i)}}(\mathbf{c})$$

$$= \rho_i \mathbf{B}^{-1}_{\overline{L(\rho_i)}}\left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N}\mathbf{p}(\mathbf{c})\right] - \mathbf{B}^{-1}_{\overline{L(\rho_i)}}\mathbf{B}_{\overline{L(\rho_i)},L(\rho_i)}\left[\mathbf{p}^{\rho_i}_{L(\rho_i)} - \mathbf{p}_{L(\rho_i)}(\mathbf{c})\right]$$

From the optimality of $\mathbf{p}(\mathbf{c})$ for the wholesale price $\mathbf{c}$, the above expression may be rewritten as:

$$= \rho_i\left[\mathbf{p}_{\overline{L(\rho_i)}}(\mathbf{c}) - \mathbf{c}_{\overline{L(\rho_i)}}\right] + \rho_i \mathbf{B}^{-1}_{\overline{L(\rho_i)}}\mathbf{B}_{\overline{L(\rho_i)},L(\rho_i)}\left[\mathbf{p}_{L(\rho_i)}(\mathbf{c}) - \mathbf{c}_{L(\rho_i)}\right]$$

$$- \mathbf{B}^{-1}_{\overline{L(\rho_i)}}\mathbf{B}_{\overline{L(\rho_i)},L(\rho_i)}\left[\mathbf{p}^{\rho_i}_{L(\rho_i)} - \mathbf{p}_{L(\rho_i)}(\mathbf{c})\right] \tag{A.22}$$

With these results in hand, we return to expressions (A.16) and (A.17), and substitute in them the inductive hypothesis and our results from expressions (A.18),(A.19), and (A.20) regarding the two parts of expression (A.17):

Hence $(\mathbf{p}^{\rho} - \mathbf{c})\left[\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c - \mathbf{B}\mathbf{p}^{\rho}\right]$

$= (\mathbf{p}^{\rho_i} - \mathbf{c})\left[\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c - \mathbf{B}\mathbf{p}^{\rho_i}\right]$

$\quad + (\mathbf{p}^{\rho} - \mathbf{c})\left[\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c - \mathbf{B}\mathbf{p}^{\rho}\right] - (\mathbf{p}^{\rho_i} - \mathbf{c})\left[\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c - \mathbf{B}\mathbf{p}^{\rho_i}\right]$

$\geq \sum_{j=1}^{i-1}(1+\rho_j)(1-\rho_j)\left[\mathbf{p}_{\widetilde{L}(\rho_j)}(\mathbf{c}) - \mathbf{c}_{\widetilde{L}(\rho_j)}\right]^t\left[\mathbf{a}_{\widetilde{L}(\rho_j)} + \mathbf{u}^c_{\widetilde{L}(\rho_j)} - \mathbf{v}^c_{\widetilde{L}(\rho_j)} - \mathbf{B}_{\widetilde{L}(\rho_j),N}\mathbf{p}(\mathbf{c})\right]$

$\quad + (1+\rho_i)(1-\rho_i)\left[\mathbf{p}_{\overline{L(\rho_{i-1})}}(\mathbf{c}) - \mathbf{c}_{\overline{L(\rho_{i-1})}}\right]^t\left[\mathbf{a}_{\overline{L(\rho_{i-1})}} + \mathbf{u}^c_{\overline{L(\rho_{i-1})}} - \mathbf{v}^c_{\overline{L(\rho_{i-1})}} - \mathbf{B}_{\overline{L(\rho_{i-1})},N}\mathbf{p}(\mathbf{c})\right]$

$\quad - (\rho - \rho_i)\left[\mathbf{p}^{\rho_i}_{L(\rho_i)} - \mathbf{c}_{L(\rho_i)}\right]^t\mathbf{B}_{L(\rho_i),\overline{L(\rho_i)}}\mathbf{B}^{-1}_{\overline{L(\rho_i)}}\left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N}\mathbf{p}(\mathbf{c})\right]$

$\quad + \left[\mathbf{p}^{\rho}_{\overline{L(\rho_i)}} - \mathbf{p}^{\rho_i}_{\overline{L(\rho_i)}}\right]^t\left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N}\mathbf{p}^{\rho}\right]$

$\quad - (\rho - \rho_i)\left[\mathbf{p}^{\rho_i}_{\overline{L(\rho_i)}} - \mathbf{c}_{\overline{L(\rho_i)}}\right]^t\left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N}\mathbf{p}(\mathbf{c})\right]$

Substituting for $\mathbf{p}^{\rho}_{\overline{L(\rho_i)}} - \mathbf{p}^{\rho_i}_{\overline{L(\rho_i)}}$ from equation (A.21) and from equation (3.9) of the optimality

conditions, we have that:

$$= \sum_{j=1}^{i-1} (1 + \rho_j)(1 - \rho_j) \left[\mathbf{p}_{\widetilde{L}(\rho_j)}(\mathbf{c}) - \mathbf{c}_{\widetilde{L}(\rho_j)} \right]^t \left[\mathbf{a}_{\widetilde{L}(\rho_j)} + \mathbf{u}^c_{\widetilde{L}(\rho_j)} - \mathbf{v}^c_{\widetilde{L}(\rho_j)} - \mathbf{B}_{\widetilde{L}(\rho_j),N} \mathbf{p}(\mathbf{c}) \right]$$

$$+ (1 + \rho_i)(1 - \rho_i) \left[\mathbf{p}_{\overline{L(\rho_{i-1})}}(\mathbf{c}) - \mathbf{c}_{\overline{L(\rho_{i-1})}} \right]^t \left[\mathbf{a}_{\overline{L(\rho_{i-1})}} + \mathbf{u}^c_{L(\rho_{i-1})} - \mathbf{v}^c_{L(\rho_{i-1})} - \mathbf{B}_{\overline{L(\rho_{i-1})},N} \mathbf{p}(\mathbf{c}) \right]$$

$$- (\rho - \rho_i) \left[\mathbf{p}^{\rho_i}_{L(\rho_i)} - \mathbf{c}_{L(\rho_i)} \right]^t \mathbf{B}_{L(\rho_i),\overline{L(\rho_i)}} \mathbf{B}^{-1}_{\overline{L(\rho_i)}} \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N} \mathbf{p}(\mathbf{c}) \right]$$

$$+ (\rho - \rho_i)(1 - \rho) \left[\mathbf{p}_{\overline{L(\rho_i)}}(\mathbf{c}) - \mathbf{c}_{\overline{L(\rho_i)}} \right]^t \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N} \mathbf{p}(\mathbf{c}) \right]$$

$$+ (\rho - \rho_i)(1 - \rho) \left[\mathbf{p}_{L(\rho_i)}(\mathbf{c}) - \mathbf{c}_{L(\rho_i)} \right]^t \mathbf{B}_{L(\rho_i),\overline{L(\rho_i)}} \mathbf{B}^{-1}_{\overline{L(\rho_i)}} \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N} \mathbf{p}(\mathbf{c}) \right]$$

$$- (\rho - \rho_i) \left[\mathbf{p}^{\rho_i}_{\overline{L(\rho_i)}} - \mathbf{p}_{\overline{L(\rho_i)}}(\mathbf{c}) \right]^t \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N} \mathbf{p}(\mathbf{c}) \right]$$

$$- (\rho - \rho_i) \left[\mathbf{p}_{\overline{L(\rho_i)}}(\mathbf{c}) - \mathbf{c}_{\overline{L(\rho_i)}} \right]^t \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N} \mathbf{p}(\mathbf{c}) \right]$$

Finally, substituting for $\left(\mathbf{p}^{\rho_i}_{\overline{L(\rho_i)}} - \mathbf{p}_{\overline{L(\rho_i)}}(\mathbf{c}) \right)$ from equation (A.22) in the above expression, we have:

$$= \sum_{j=1}^{i-1} (1 + \rho_j)(1 - \rho_j) \left[\mathbf{p}_{\widetilde{L}(\rho_j)}(\mathbf{c}) - \mathbf{c}_{\widetilde{L}(\rho_j)} \right]^t \left[\mathbf{a}_{\widetilde{L}(\rho_j)} + \mathbf{u}^c_{\widetilde{L}(\rho_j)} - \mathbf{v}^c_{\widetilde{L}(\rho_j)} - \mathbf{B}_{\widetilde{L}(\rho_j),N} \mathbf{p}(\mathbf{c}) \right]$$

$$+ (1 + \rho_i)(1 - \rho_i) \left[\mathbf{p}_{\overline{L(\rho_{i-1})}}(\mathbf{c}) - \mathbf{c}_{\overline{L(\rho_{i-1})}} \right]^t \left[\mathbf{a}_{\overline{L(\rho_{i-1})}} + \mathbf{u}^c_{L(\rho_{i-1})} - \mathbf{v}^c_{L(\rho_{i-1})} - \mathbf{B}_{\overline{L(\rho_{i-1})},N} \mathbf{p}(\mathbf{c}) \right]$$

$$- (\rho - \rho_i) \left[\mathbf{p}^{\rho_i}_{L(\rho_i)} - \mathbf{c}_{L(\rho_i)} \right]^t \mathbf{B}_{L(\rho_i),\overline{L(\rho_i)}} \mathbf{B}^{-1}_{\overline{L(\rho_i)}} \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N} \mathbf{p}(\mathbf{c}) \right]$$

$$+ (\rho - \rho_i)(1 - \rho) \left[\mathbf{p}_{\overline{L(\rho_i)}}(\mathbf{c}) - \mathbf{c}_{\overline{L(\rho_i)}} \right]^t \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N} \mathbf{p}(\mathbf{c}) \right]$$

$$+ (\rho - \rho_i)(1 - \rho) \left[\mathbf{p}_{L(\rho_i)}(\mathbf{c}) - \mathbf{c}_{L(\rho_i)} \right]^t \mathbf{B}_{L(\rho_i),\overline{L(\rho_i)}} \mathbf{B}^{-1}_{\overline{L(\rho_i)}} \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N} \mathbf{p}(\mathbf{c}) \right]$$

$$- (\rho - \rho_i)(1 + \rho_i) \left[\mathbf{p}_{\overline{L(\rho_i)}}(\mathbf{c}) - \mathbf{c}_{\overline{L(\rho_i)}} \right]^t \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N} \mathbf{p}(\mathbf{c}) \right]$$

$$- (\rho - \rho_i) \rho_i \left[\mathbf{p}_{L(\rho_i)}(\mathbf{c}) - \mathbf{c}_{L(\rho_i)} \right]^t \mathbf{B}_{L(\rho_i),\overline{L(\rho_i)}} \mathbf{B}^{-1}_{\overline{L(\rho_i)}} \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N} \mathbf{p}(\mathbf{c}) \right]$$

$$+ (\rho - \rho_i) \left[\mathbf{p}^{\rho_i}_{L(\rho_i)} - \mathbf{p}_{L(\rho_i)}(\mathbf{c}) \right]^t \mathbf{B}_{L(\rho_i),\overline{L(\rho_i)}} \mathbf{B}^{-1}_{\overline{L(\rho_i)}} \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N} \mathbf{p}(\mathbf{c}) \right]$$

Via simple rearrangements and carefully observing the signs of expressions, this reduces to:

$$\geq \sum_{j=1}^{i} (1 + \rho_j)(1 - \rho_j) \left[\mathbf{p}_{\widetilde{L}(\rho_j)}(\mathbf{c}) - \mathbf{c}_{\widetilde{L}(\rho_j)} \right]^t \left[\mathbf{a}_{\widetilde{L}(\rho_j)} + \mathbf{u}^c_{\widetilde{L}(\rho_j)} - \mathbf{v}^c_{\widetilde{L}(\rho_j)} - \mathbf{B}_{\widetilde{L}(\rho_j),N} \mathbf{p}(\mathbf{c}) \right]$$

$$+ (1 + \rho_i)(1 - \rho_i) \left[\mathbf{p}_{\overline{L(\rho_i)}}(\mathbf{c}) - \mathbf{c}_{\overline{L(\rho_i)}} \right]^t \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N} \mathbf{p}(\mathbf{c}) \right]$$

$$- (\rho - \rho_i) \left[\mathbf{p}^{\rho_i}_{L(\rho_i)} - \mathbf{c}_{L(\rho_i)} \right]^t \mathbf{B}_{L(\rho_i),\overline{L(\rho_i)}} \mathbf{B}^{-1}_{\overline{L(\rho_i)}} \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N} \mathbf{p}(\mathbf{c}) \right]$$

$$+ (\rho - \rho_i)(1 - \rho) \left[\mathbf{p}_{\overline{L(\rho_i)}}(\mathbf{c}) - \mathbf{c}_{\overline{L(\rho_i)}} \right]^t \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N} \mathbf{p}(\mathbf{c}) \right]$$

$$+ (\rho - \rho_i) \left[\mathbf{p}^{\rho_i}_{L(\rho_i)} - \mathbf{c}_{L(\rho_i)} \right]^t \mathbf{B}_{L(\rho_i),\overline{L(\rho_i)}} \mathbf{B}^{-1}_{\overline{L(\rho_i)}} \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N} \mathbf{p}(\mathbf{c}) \right]$$

$$- (\rho - \rho_i)(1 + \rho_i) \left[\mathbf{p}_{\overline{L(\rho_i)}}(\mathbf{c}) - \mathbf{c}_{\overline{L(\rho_i)}} \right]^t \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N} \mathbf{p}(\mathbf{c}) \right]$$

$$- \left(\rho - \rho_i\right) \rho_i \left[\mathbf{p}_{L(\rho_i)}\left(\mathbf{c}\right) - \mathbf{c}_{L(\rho_i)}\right]^t \mathbf{B}_{L(\rho_i),\overline{L(\rho_i)}} \mathbf{B}^{-1}_{\overline{L(\rho_i)}} \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N}\mathbf{p}\left(\mathbf{c}\right)\right]$$

Collecting common coefficients:

$$= \sum_{j=1}^i \left(1 + \rho_j\right)\left(1 - \rho_j\right) \left[\mathbf{p}_{\widetilde{L}(\rho_j)}\left(\mathbf{c}\right) - \mathbf{c}_{\widetilde{L}(\rho_j)}\right]^t \left[\mathbf{a}_{\widetilde{L}(\rho_j)} + \mathbf{u}^c_{\widetilde{L}(\rho_j)} - \mathbf{v}^c_{\widetilde{L}(\rho_j)} - \mathbf{B}_{\widetilde{L}(\rho_j),N}\mathbf{p}\left(\mathbf{c}\right)\right]$$

$$+ \left(1 + \rho\right)\left(1 - \rho\right) \left[\mathbf{p}_{\overline{L(\rho_i)}}\left(\mathbf{c}\right) - \mathbf{c}_{\overline{L(\rho_i)}}\right]^t \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N}\mathbf{p}\left(\mathbf{c}\right)\right]$$

$$- \left(\rho - \rho_i\right) \rho_i \left[\mathbf{p}_{L(\rho_i)}\left(\mathbf{c}\right) - \mathbf{c}_{L(\rho_i)}\right]^t \mathbf{B}_{L(\rho_i),\overline{L(\rho_i)}} \mathbf{B}^{-1}_{\overline{L(\rho_i)}} \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N}\mathbf{p}\left(\mathbf{c}\right)\right]$$

$$\geq \sum_{j=1}^i \left(1 + \rho_j\right)\left(1 - \rho_j\right) \left[\mathbf{p}_{\widetilde{L}(\rho_j)}\left(\mathbf{c}\right) - \mathbf{c}_{\widetilde{L}(\rho_j)}\right]^t \left[\mathbf{a}_{\widetilde{L}(\rho_j)} + \mathbf{u}^c_{\widetilde{L}(\rho_j)} - \mathbf{v}^c_{\widetilde{L}(\rho_j)} - \mathbf{B}_{\widetilde{L}(\rho_j),N}\mathbf{p}\left(\mathbf{c}\right)\right]$$

$$+ \left(1 + \rho\right)\left(1 - \rho\right) \left[\mathbf{p}_{\overline{L(\rho_i)}}\left(\mathbf{c}\right) - \mathbf{c}_{\overline{L(\rho_i)}}\right]^t \left[\mathbf{a}_{\overline{L(\rho_i)}} + \mathbf{u}^c_{\overline{L(\rho_i)}} - \mathbf{v}^c_{\overline{L(\rho_i)}} - \mathbf{B}_{\overline{L(\rho_i)},N}\mathbf{p}\left(\mathbf{c}\right)\right]. \qquad \blacksquare$$

A.0.5 Proof of Lemma 3.21

Proof. Observe that:

$$\mathbf{a}_{\widetilde{\varsigma(1/2)}} + \mathbf{u}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{v}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{B}_{\widetilde{\varsigma(1/2)},N}\widetilde{\mathbf{p}}^{1/2}$$

$$= \mathbf{a}^c_{\widetilde{\varsigma(1/2)}} + \mathbf{u}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{v}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{B}_{\widetilde{\varsigma(1/2)},\overline{G}}\widetilde{\mathbf{p}}^{1/2}_{\overline{G}}$$

by absorbing terms into $\mathbf{a}^c$. Substituting the optimality condition of equation (3.14) for $\rho = 1/2$, we have further that:

$$= \tfrac{1}{2}\left(\mathbf{a}^c_{\widetilde{\varsigma(1/2)}} + \mathbf{u}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{v}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{B}_{\widetilde{\varsigma(1/2)},\overline{G}}\mathbf{p}^G_{\overline{G}}\right)$$

$$= \tfrac{1}{2}\left(\mathbf{a}_{\widetilde{\varsigma(1/2)}} + \mathbf{u}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{v}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{B}_{\widetilde{\varsigma(1/2)},N}\mathbf{p}^G\right)$$

Rearranging the above equation, we now have that:

$$\mathbf{B}_{\widetilde{\varsigma(1/2)}}\widetilde{\mathbf{p}}^{1/2}_{\widetilde{\varsigma(1/2)}} = \frac{1}{2}\left(\mathbf{a}_{\widetilde{\varsigma(1/2)}} + \mathbf{u}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{v}^c_{\widetilde{\varsigma(1/2)}}\right) + \frac{1}{2}\mathbf{B}_{\widetilde{\varsigma(1/2)},N}\mathbf{p}^G$$

$$- \mathbf{B}_{\widetilde{\varsigma(1/2)},G}\mathbf{p}^G_G - \mathbf{B}_{\widetilde{\varsigma(1/2)},\varsigma(1/2)}\mathbf{p}^u_{\varsigma(1/2)}. \qquad (A.23)$$

Since $\mathbf{p}\left(\mathbf{c}\right) \leq \mathbf{p}^G$ (Lemma 3.13), the optimality conditions for $\mathbf{p}^G$ and $\mathbf{p(c)}$ give us that:

$$\mathbf{a}_{\widetilde{\varsigma(1/2)}} + \mathbf{u}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{v}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{B}_{\widetilde{\varsigma(1/2)},N}\mathbf{p}\left(\mathbf{c}\right) - \mathbf{B}_{\widetilde{\varsigma(1/2)},N}\left[\mathbf{p}\left(\mathbf{c}\right) - \mathbf{c}\right] = 0 \text{ and}$$

$$\mathbf{a}_{\widetilde{\varsigma(1/2)}} + \mathbf{u}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{v}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{B}_{\widetilde{\varsigma(1/2)},N}\mathbf{p}^G - \mathbf{B}_{\widetilde{\varsigma(1/2)},N}\left[\mathbf{p}^G - \mathbf{c}\right] = 0$$

The two statements above imply that $\mathbf{B}_{\widetilde{\varsigma(1/2)},N}\mathbf{p}^G = \mathbf{B}_{\widetilde{\varsigma(1/2)},N}\mathbf{p}\left(\mathbf{c}\right)$.

Now, from equation (A.23), $\widetilde{\mathbf{p}}^{1/2}_{\varsigma(1/2)}$ may be evaluated to be:

$$\mathbf{p}^u_{\varsigma\widetilde{(1/2)}} > \widetilde{\mathbf{p}}^{1/2}_{\varsigma\widetilde{(1/2)}}$$

$$= \tfrac{1}{2}\mathbf{B}^{-1}_{\varsigma\widetilde{(1/2)}}\left(\mathbf{a}_{\varsigma\widetilde{(1/2)}} + \mathbf{u}^c_{\varsigma\widetilde{(1/2)}} - \mathbf{v}^c_{\varsigma\widetilde{(1/2)}}\right) + \tfrac{1}{2}\mathbf{B}^{-1}_{\varsigma\widetilde{(1/2)}}\mathbf{B}_{\varsigma\widetilde{(1/2)},N}\mathbf{p}^G - \mathbf{B}^{-1}_{\varsigma\widetilde{(1/2)}}\mathbf{B}_{\varsigma\widetilde{(1/2)},G}\mathbf{p}^G_G$$

$$-\mathbf{B}^{-1}_{\varsigma\widetilde{(1/2)}}\mathbf{B}_{\varsigma\widetilde{(1/2)},\varsigma(1/2)}\mathbf{p}^u_{\varsigma(1/2)}$$

Using the fact that $\mathbf{B}_{\varsigma\widetilde{(1/2)},N}\mathbf{p}^G = \mathbf{B}_{\varsigma\widetilde{(1/2)},N}\mathbf{p}\,(\mathbf{c})$, we further infer that:

$$= \tfrac{1}{2}\mathbf{B}^{-1}_{\varsigma\widetilde{(1/2)}}\left(\mathbf{a}_{\varsigma\widetilde{(1/2)}} + \mathbf{u}^c_{\varsigma\widetilde{(1/2)}} - \mathbf{v}^c_{\varsigma\widetilde{(1/2)}}\right) + \tfrac{1}{2}\mathbf{B}^{-1}_{\varsigma\widetilde{(1/2)}}\mathbf{B}_{\varsigma\widetilde{(1/2)},N}\mathbf{p}\,(\mathbf{c}) - \mathbf{B}^{-1}_{\varsigma\widetilde{(1/2)}}\mathbf{B}_{\varsigma\widetilde{(1/2)},G}\mathbf{p}^G_G$$

$$-\mathbf{B}^{-1}_{\varsigma\widetilde{(1/2)}}\mathbf{B}_{\varsigma\widetilde{(1/2)},\varsigma(1/2)}\mathbf{p}^u_{\varsigma(1/2)}$$

$$\geq \tfrac{1}{2}\mathbf{B}^{-1}_{\varsigma\widetilde{(1/2)}}\left(\mathbf{a}_{\varsigma\widetilde{(1/2)}} + \mathbf{u}^c_{\varsigma\widetilde{(1/2)}} - \mathbf{v}^c_{\varsigma\widetilde{(1/2)}}\right) + \tfrac{1}{2}\mathbf{B}^{-1}_{\varsigma\widetilde{(1/2)}}\mathbf{B}_{\varsigma\widetilde{(1/2)},N}\mathbf{p}\,(\mathbf{c}) - \mathbf{B}^{-1}_{\varsigma\widetilde{(1/2)}}\mathbf{B}_{\varsigma\widetilde{(1/2)},G}\mathbf{p}^{1/2}_G$$

$$-\mathbf{B}^{-1}_{\varsigma\widetilde{(1/2)}}\mathbf{B}_{\varsigma\widetilde{(1/2)},\varsigma(1/2)}\mathbf{p}^u_{\varsigma(1/2)} \quad = \quad \mathbf{p}^{1/2}_{\varsigma\widetilde{(1/2)}}$$

Hence $\mathbf{p}^u_{\varsigma\widetilde{(1/2)}} > \widetilde{\mathbf{p}}^{1/2}_{\varsigma\widetilde{(1/2)}} \geq \mathbf{p}^{1/2}_{\varsigma\widetilde{(1/2)}}$. Furthermore, from the definition of G, we have that $\mathbf{p}^{1/2}_G < \mathbf{p}^c_G \leq \mathbf{p}^u_G$. It then follows that $L(1/2) \subseteq \varsigma(1/2)$, and hence $\mathbf{p}^{1/2} \leq \widetilde{\mathbf{p}}^{1/2}$. ∎

A.0.6 Proof of Lemma 3.22

Proof. The proof is via comparisons of $\mathbf{w}^{HC}_G$ with $\mathbf{w}^{1/2}_G$ through a couple of wholesale price vectors corresponding to combinations of $\widetilde{\mathbf{p}}^{1/2}$ and $\mathbf{p}^{1/2}$. More specifically, we construct two new wholesale price vectors $\overline{\mathbf{w}}$ and $\widetilde{\mathbf{w}}$ and by establishing key properties are relations between the two, we prove the result.

Consider $\overline{\mathbf{p}}_G = \mathbf{p}^c_G$ and $\overline{\mathbf{p}}_{\overline{G}} = \mathbf{p}^{1/2}_{\overline{G}}$. We observe that $G \subseteq \overline{L(1/2)}$ by the definition of G. Let $\overline{\mathbf{w}}$ be defined as follows:

$$\overline{\mathbf{w}}_{L(1/2)} = \mathbf{p}^u_{L(1/2)} \text{ and}$$

$$\overline{\mathbf{w}}_{\overline{L(1/2)}} = 2\overline{\mathbf{p}}_{\overline{L(1/2)}} - \mathbf{B}^{-1}_{\overline{L(1/2)}}\left[\mathbf{a}_{\overline{L(1/2)}} + \mathbf{u}^c_{\overline{L(1/2)}} - \mathbf{v}^c_{\overline{L(1/2)}} - \mathbf{B}_{\overline{L(1/2)},L(1/2)}\mathbf{p}^u_{L(1/2)}\right].$$

Via an argument similar to that of Lemma 3.19, one may verify that $\mathbf{p}(\overline{\mathbf{w}}) = \overline{\mathbf{p}}$. With a little more care, it may also be observed that $\overline{\mathbf{w}}_{\overline{G}} = \mathbf{w}^{1/2}_{\overline{G}}$.

From the optimality condition of equation (3.14) for $\rho = 1/2$, we have that:

$$\mathbf{a}_G + \mathbf{u}^c_G - \mathbf{v}^c_G - \mathbf{B}_{G,N}\mathbf{p}^{1/2} = \frac{1}{2}\left[\mathbf{a}_G + \mathbf{u}^c_G - \mathbf{v}^c_G - \mathbf{B}_{G,N}\mathbf{p}\,(\mathbf{c})\right] \qquad (A.24)$$

155

and from the optimality of $\mathbf{p}(\mathbf{c})$ for a wholesale price of $\mathbf{c}$ and definition of G, we have that:

$$\mathbf{a}_G + \mathbf{u}_G^c - \mathbf{v}_G^c - \mathbf{B}_{G,N}\mathbf{p}(\mathbf{c}) = \mathbf{B}_{G,N}\left[\mathbf{p}(\mathbf{c}) - \mathbf{c}\right] \tag{A.25}$$

Rearranging equation (A.24), we may obtain $\mathbf{p}_G^{1/2}$ to be:

$\mathbf{p}_G^c > \mathbf{p}_G^{1/2}$

$= \frac{1}{2}\mathbf{B}_G^{-1}\left(\mathbf{a}_G + \mathbf{u}_G^c - \mathbf{v}_G^c\right) + \frac{1}{2}\mathbf{B}_G^{-1}\mathbf{B}_{G,N}\mathbf{p}(\mathbf{c}) - \mathbf{B}_G^{-1}\mathbf{B}_{G,\overline{G}}\mathbf{p}_{\overline{G}}^{1/2}$

We substitute for the value of $\mathbf{p}_G(\mathbf{c})$ from equation (A.25) in the above expression, which yields

$= \frac{1}{2}\mathbf{B}_G^{-1}\left(\mathbf{a}_G + \mathbf{u}_G^c - \mathbf{v}_G^c\right) + \frac{1}{2}\mathbf{B}_G^{-1}\mathbf{B}_{G,\overline{G}}\mathbf{p}_{\overline{G}}(\mathbf{c}) - \mathbf{B}_G^{-1}\mathbf{B}_{G,\overline{G}}\mathbf{p}_{\overline{G}}^{1/2}$

$\qquad + \frac{1}{2}\left[\frac{1}{2}\mathbf{B}_G^{-1}\left(\mathbf{a}_G + \mathbf{u}_G^c - \mathbf{v}_G^c\right) + \frac{1}{2}\mathbf{B}_G^{-1}\mathbf{B}_{G,N}\mathbf{c} - \mathbf{B}_G^{-1}\mathbf{B}_{G,\overline{G}}\mathbf{p}_{\overline{G}}(\mathbf{c})\right]$

$= \frac{3}{4}\mathbf{B}_G^{-1}\left(\mathbf{a}_G + \mathbf{u}_G^c - \mathbf{v}_G^c\right) + \frac{1}{4}\mathbf{B}_G^{-1}\mathbf{B}_{G,N}\mathbf{c} - \mathbf{B}_G^{-1}\mathbf{B}_{G,\overline{G}}\mathbf{p}_{\overline{G}}^{1/2}.$

From the optimality of $\overline{\mathbf{p}}$ for the wholesale price $\overline{\mathbf{w}}$ we have that:

$$\mathbf{a}_G + \mathbf{u}_G^c - \mathbf{v}_G^c - \mathbf{B}_{G,N}\overline{\mathbf{p}} = \mathbf{B}_{G,N}\left(\overline{\mathbf{p}} - \overline{\mathbf{w}}\right) \tag{A.26}$$

Following equation (A.26), we may write $\overline{\mathbf{w}}_G$ as:

$\overline{\mathbf{w}}_G = 2\mathbf{p}_G^c - \mathbf{B}_G^{-1}\left(\mathbf{a}_G + \mathbf{u}_G^c - \mathbf{v}_G^c\right) + 2\mathbf{B}_G^{-1}\mathbf{B}_{G,\overline{G}}\overline{\mathbf{p}}_{\overline{G}} - \mathbf{B}_G^{-1}\mathbf{B}_{G,\overline{G}}\overline{\mathbf{w}}_{\overline{G}}$

Substituting for $\frac{3}{2}\mathbf{p}_G^c$ from the previous result for $\mathbf{p}_G^c$:

$> \frac{1}{2}\mathbf{p}_G^c + \frac{9}{8}\mathbf{B}_G^{-1}\left(\mathbf{a}_G + \mathbf{u}_G^c - \mathbf{v}_G^c\right) + \frac{3}{8}\mathbf{B}_G^{-1}\mathbf{B}_{G,N}\mathbf{c} - \frac{3}{2}\mathbf{B}_G^{-1}\mathbf{B}_{G,\overline{G}}\mathbf{p}_{\overline{G}}^{1/2}$

$\quad - \mathbf{B}_G^{-1}\left(\mathbf{a}_G + \mathbf{u}_G^c - \mathbf{v}_G^c\right) + 2\mathbf{B}_G^{-1}\mathbf{B}_{G,\overline{G}}\overline{\mathbf{p}}_{\overline{G}} - \mathbf{B}_G^{-1}\mathbf{B}_{G,\overline{G}}\overline{\mathbf{w}}_{\overline{G}}$

$= \frac{1}{2}\mathbf{p}_G^c + \frac{1}{8}\mathbf{B}_G^{-1}\left(\mathbf{a}_G + \mathbf{u}_G^c - \mathbf{v}_G^c\right) + \frac{3}{8}\mathbf{B}_G^{-1}\mathbf{B}_{G,N}\mathbf{c} + \frac{1}{2}\mathbf{B}_G^{-1}\mathbf{B}_{G,\overline{G}}\mathbf{p}_{\overline{G}}^{1/2} - \mathbf{B}_G^{-1}\mathbf{B}_{G,\overline{G}}\overline{\mathbf{w}}_{\overline{G}}$

Since $\mathbf{a}_G \geq \mathbf{B}_{G,N}\mathbf{c}$:

$\geq \frac{1}{2}\mathbf{p}_G^c + \frac{1}{2}\mathbf{B}_G^{-1}\mathbf{B}_{G,N}\mathbf{c} + \frac{1}{2}\mathbf{B}_G^{-1}\mathbf{B}_{G,\overline{G}}\mathbf{p}_{\overline{G}}^{1/2} - \mathbf{B}_G^{-1}\mathbf{B}_{G,\overline{G}}\overline{\mathbf{w}}_{\overline{G}}$

$= \frac{1}{2}\mathbf{p}_G^c + \frac{1}{2}\mathbf{c}_G + \frac{1}{2}\mathbf{B}_G^{-1}\mathbf{B}_{G,\overline{G}}\mathbf{c}_{\overline{G}} + \frac{1}{2}\mathbf{B}_G^{-1}\mathbf{B}_{G,\overline{G}}\mathbf{p}_{\overline{G}}^{1/2} - \mathbf{B}_G^{-1}\mathbf{B}_{G,\overline{G}}\mathbf{w}_{\overline{G}}^{1/2}.$

By the definition of $\mathbf{w}_{\overline{L(1/2)}}^{1/2}$, we have that:

$$\mathbf{w}^{1/2}_{\overline{L(1/2)}} = 2\mathbf{p}^{1/2}_{\overline{L(1/2)}} - \mathbf{B}^{-1}_{\overline{L(1/2)}}\left[\mathbf{a}_{\overline{L(1/2)}} + \mathbf{u}^c_{\overline{L(1/2)}} - \mathbf{v}^c_{\overline{L(1/2)}} - \mathbf{B}_{\overline{L(1/2)},L(1/2)}\mathbf{p}^u_{L(1/2)}\right]$$

with in turns implies that:

$$\mathbf{p}^{1/2}_{\overline{L(1/2)}} - \mathbf{w}^{1/2}_{\overline{L(1/2)}} = \mathbf{B}^{-1}_{\overline{L(1/2)}}\left[\mathbf{a}_{\overline{L(1/2)}} + \mathbf{u}^c_{\overline{L(1/2)}} - \mathbf{v}^c_{\overline{L(1/2)}} - \mathbf{B}_{\overline{L(1/2)},N}\mathbf{p}^{1/2}\right]$$

$$= \tfrac{1}{2}\mathbf{B}^{-1}_{\overline{L(1/2)}}\left[\mathbf{a}_{\overline{L(1/2)}} + \mathbf{u}^c_{\overline{L(1/2)}} - \mathbf{v}^c_{\overline{L(1/2)}} - \mathbf{B}_{\overline{L(1/2)},N}\mathbf{p}\,(\mathbf{c})\right] \qquad \text{(from equation (3.9))}$$

$$= \tfrac{1}{2}\mathbf{B}^{-1}_{\overline{L(1/2)}}\mathbf{B}_{\overline{L(1/2)},N}\left[\mathbf{p}\,(\mathbf{c}) - \mathbf{c}\right] \qquad \text{(from optimality of } \mathbf{p}(\mathbf{c}) \text{ for wholesale price } \mathbf{c})$$

$$= \tfrac{1}{2}\left[\mathbf{p}_{\overline{L(1/2)}}\,(\mathbf{c}) - \mathbf{c}_{\overline{L(1/2)}}\right] + \tfrac{1}{2}\mathbf{B}^{-1}_{\overline{L(1/2)}}\mathbf{B}_{\overline{L(1/2)},L(1/2)}\left[\mathbf{p}_{L(1/2)}\,(\mathbf{c}) - \mathbf{c}_{L(1/2)}\right]$$

$$\leq \tfrac{1}{2}\left[\mathbf{p}_{\overline{L(1/2)}}\,(\mathbf{c}) - \mathbf{c}_{\overline{L(1/2)}}\right] \qquad\qquad \text{. (since } \mathbf{p}(\mathbf{c}) \geq \mathbf{c})$$

If in the definition of $\mathbf{w}^{1/2}_{\overline{L(1/2)}}$ noted earlier in the proof, we plug in the formula for $\mathbf{p}^{1/2}_{\overline{L(1/2)}}$, through straightforward algebraic manipulations, we obtain that:

$$\mathbf{w}^{1/2}_{\overline{L(1/2)}} = \mathbf{p}_{\overline{L(1/2)}}\,(\mathbf{c}) - \mathbf{B}^{-1}_{\overline{L(1/2)}}\mathbf{B}_{\overline{L(1/2)},L(1/2)}\left(\mathbf{p}^u_{L(1/2)} - \mathbf{p}_{L(1/2)}(\mathbf{c})\right)$$

from which we may infer that $\mathbf{w}^{1/2}_{\overline{L(1/2)}} \geq \mathbf{p}_{\overline{L(1/2)}}\,(\mathbf{c})$. As a consequence of the previous fact and this one, it then follows that:

$$\mathbf{w}^{1/2}_{\overline{L(1/2)}} - \mathbf{c}_{\overline{L(1/2)}} \geq \mathbf{p}_{\overline{L(1/2)}}\,(\mathbf{c}) - \mathbf{c}_{\overline{L(1/2)}} \geq 2\left(\mathbf{p}^{1/2}_{\overline{L(1/2)}} - \mathbf{w}^{1/2}_{\overline{L(1/2)}}\right). \qquad (A.27)$$

Armed with these results, we return our discussion following equation (A.26) regarding a lower bound on $\overline{\mathbf{w}}_G$:

$$\overline{\mathbf{w}}_G \geq \tfrac{1}{2}\mathbf{p}^c_G + \tfrac{1}{2}\mathbf{c}_G + \tfrac{1}{2}\mathbf{B}^{-1}_G\mathbf{B}_{G,\overline{G}}\mathbf{c}_{\overline{G}} + \tfrac{1}{2}\mathbf{B}^{-1}_G\mathbf{B}_{G,\overline{G}}\mathbf{p}^{1/2}_{\overline{G}} - \mathbf{B}^{-1}_G\mathbf{B}_{G,\overline{G}}\mathbf{w}^{1/2}_{\overline{G}}$$

$$= \tfrac{1}{2}\mathbf{p}^c_G + \tfrac{1}{2}\mathbf{c}_G + \tfrac{1}{2}\mathbf{B}^{-1}_G\mathbf{B}_{G,\overline{G}\cap L(1/2)}\mathbf{c}_{\overline{G}\cap L(1/2)} + \tfrac{1}{2}\mathbf{B}^{-1}_G\mathbf{B}_{G,\overline{G}\cap L(1/2)}\mathbf{p}^{1/2}_{\overline{G}\cap L(1/2)}$$

$$\quad -\mathbf{B}^{-1}_G\mathbf{B}_{G,\overline{G}\cap L(1/2)}\mathbf{w}^{1/2}_{\overline{G}\cap L(1/2)} + \tfrac{1}{2}\mathbf{B}^{-1}_G\mathbf{B}_{G,\overline{G}\cap\overline{L(1/2)}}\mathbf{c}_{\overline{G}\cap\overline{L(1/2)}} + \tfrac{1}{2}\mathbf{B}^{-1}_G\mathbf{B}_{G,\overline{G}\cap\overline{L(1/2)}}\mathbf{p}^{1/2}_{\overline{G}\cap\overline{L(1/2)}}$$

$$\quad -\mathbf{B}^{-1}_G\mathbf{B}_{G,\overline{G}\cap\overline{L(1/2)}}\mathbf{w}^{1/2}_{\overline{G}\cap\overline{L(1/2)}}$$

Noting that $\mathbf{w}^{1/2}_{L(1/2)} = \mathbf{p}^u_{L(1/2)} = \mathbf{p}^{1/2}_{L(1/2)}$, the above expression may be simplified as:

$$= \tfrac{1}{2}\mathbf{p}^c_G + \tfrac{1}{2}\mathbf{c}_G - \tfrac{1}{2}\mathbf{B}^{-1}_G\mathbf{B}_{G,\overline{G}\cap L(1/2)}\left(\mathbf{p}^u_{\overline{G}\cap L(1/2)} - \mathbf{c}_{\overline{G}\cap L(1/2)}\right)$$

$$\quad -\tfrac{1}{2}\mathbf{B}^{-1}_G\mathbf{B}_{G,\overline{G}\cap\overline{L(1/2)}}\left[\left(\mathbf{w}^{1/2}_{\overline{G}\cap\overline{L(1/2)}} - \mathbf{c}_{\overline{G}\cap\overline{L(1/2)}}\right) - \left(\mathbf{p}^{1/2}_{\overline{G}\cap\overline{L(1/2)}} - \mathbf{w}^{1/2}_{\overline{G}\cap\overline{L(1/2)}}\right)\right]$$

$$\geq \tfrac{1}{2}\mathbf{p}^c_G + \tfrac{1}{2}\mathbf{c}_G,$$

as a consequence of equation (A.27) and straightforward sign arguments about $\mathbf{B}$ and $\mathbf{B}^{-1}$.

From a discussion prior to this proof, we know that $\mathbf{p}(\mathbf{w}^{HC}) = \widetilde{\mathbf{p}}^{1/2}$. To complete this proof, we consider an alternate $\widetilde{\mathbf{w}}$ that would also induce a retail price of $\widetilde{\mathbf{p}}^{1/2}$. Suppose:

$$\widetilde{\mathbf{w}}_{L(1/2)} = \mathbf{p}^u_{L(1/2)} \text{ and}$$

$$\widetilde{\mathbf{w}}_{\overline{L(1/2)}} = 2\widetilde{\mathbf{p}}_{\overline{L(1/2)}}^{1/2} - \mathbf{B}_{\overline{L(1/2)}}^{-1} \left[\mathbf{a}_{\overline{L(1/2)}} + \mathbf{u}_{\overline{L(1/2)}}^c - \mathbf{v}_{\overline{L(1/2)}}^c - \mathbf{B}_{\overline{L(1/2)},L(1/2)} \mathbf{p}_{L(1/2)}^u \right].$$

It is not hard to show that $\mathbf{p}(\widetilde{\mathbf{w}}) = \widetilde{\mathbf{p}}^{1/2}$. Furthermore, from the definitions of $\widetilde{\mathbf{w}}_{\overline{L(1/2)}}$ and $\overline{\mathbf{w}}_{\overline{L(1/2)}}$ it follows easily that:

$$2\widetilde{\mathbf{p}}_{\overline{L(1/2)}}^{1/2} - \widetilde{\mathbf{w}}_{\overline{L(1/2)}} = 2\overline{\mathbf{p}}_{\overline{L(1/2)}} - \overline{\mathbf{w}}_{\overline{L(1/2)}}$$

$$\Rightarrow \widetilde{\mathbf{p}}_{\overline{L(1/2)}}^{1/2} - \widetilde{\mathbf{w}}_{\overline{L(1/2)}} = \overline{\mathbf{p}}_{\overline{L(1/2)}} - \overline{\mathbf{w}}_{\overline{L(1/2)}} - \left(\widetilde{\mathbf{p}}_{\overline{L(1/2)}}^{1/2} - \overline{\mathbf{p}}_{\overline{L(1/2)}} \right) \le \overline{\mathbf{p}}_{\overline{L(1/2)}} - \overline{\mathbf{w}}_{\overline{L(1/2)}}. \quad (A.28)$$

The last inequality in (A.28) may be inferred from the definition of $\overline{\mathbf{p}}$ and the result that $\widetilde{\mathbf{p}}^{1/2} \ge \mathbf{p}^{1/2}$ (Lemma 3.21). Moreover $\widetilde{\mathbf{p}}_G^{1/2} = \mathbf{p}_G^c = \overline{\mathbf{p}}_G$ implies in (A.28) that $\widetilde{\mathbf{p}}_G^{1/2} - \widetilde{\mathbf{w}}_G = \overline{\mathbf{p}}_G - \overline{\mathbf{w}}_G$ and hence, $\widetilde{\mathbf{w}}_G = \overline{\mathbf{w}}_G$.

From the definition of $\widetilde{\mathbf{w}}_{\overline{\tau}}^{HC}$, we may infer that:

$$\mathbf{B}_{\overline{\tau}} \left(\widetilde{\mathbf{p}}_{\overline{\tau}}^{1/2} - \mathbf{w}_{\overline{\tau}}^{HC} \right) = \mathbf{a}_{\overline{\tau}} + \mathbf{u}_{\overline{\tau}}^c - \mathbf{v}_{\overline{\tau}}^c - \mathbf{B}_{\overline{\tau},N} \widetilde{\mathbf{p}}^{1/2} \quad (A.29)$$

From Lemma 3.21, we have that $\widetilde{\mathbf{p}}^{1/2} \ge \mathbf{p}^{1/2}$. This would imply that $L(1/2) \subseteq \tau$ (recall that $\tau = \{i \in N : \widetilde{p}_i^{1/2} = p_i^u\}$) or that $\overline{\tau} \subseteq \overline{L(1/2)}$. Using this fact in the definition of $\widetilde{\mathbf{w}}_{\overline{L(1/2)}}$, we then have that:

$$\mathbf{B}_{\overline{\tau},\overline{L(1/2)}} \left(\widetilde{\mathbf{p}}_{\overline{L(1/2)}}^{1/2} - \widetilde{\mathbf{w}}_{\overline{L(1/2)}} \right) = \mathbf{a}_{\overline{\tau}} + \mathbf{u}_{\overline{\tau}}^c - \mathbf{v}_{\overline{\tau}}^c - \mathbf{B}_{\overline{\tau},N} \widetilde{\mathbf{p}}^{1/2} \quad (A.30)$$

From equations (A.29) and (A.30), we can infer that:

$$\widetilde{\mathbf{p}}_{\overline{\tau}}^{1/2} - \mathbf{w}_{\overline{\tau}}^{HC} = \mathbf{B}_{\overline{\tau}}^{-1} \mathbf{B}_{\overline{\tau},\overline{L(1/2)}} \left(\widetilde{\mathbf{p}}_{\overline{L(1/2)}}^{1/2} - \widetilde{\mathbf{w}}_{\overline{L(1/2)}} \right)$$

$$= \widetilde{\mathbf{p}}_{\overline{\tau}}^{1/2} - \widetilde{\mathbf{w}}_{\overline{\tau}} + \mathbf{B}_{\overline{\tau}}^{-1} \mathbf{B}_{\overline{\tau},\overline{L(1/2)}\backslash\overline{\tau}} \left(\widetilde{\mathbf{p}}_{\overline{L(1/2)}\backslash\overline{\tau}}^{1/2} - \widetilde{\mathbf{w}}_{\overline{L(1/2)}\backslash\overline{\tau}} \right) \le \widetilde{\mathbf{p}}_{\overline{\tau}}^{1/2} - \widetilde{\mathbf{w}}_{\overline{\tau}}$$

where the last inequality follows from observing that $\mathbf{p}(\widetilde{\mathbf{w}}) = \widetilde{\mathbf{p}}^{1/2}$. More importantly, we may conclude that $\mathbf{w}_{\overline{\tau}}^{HC} \ge \widetilde{\mathbf{w}}_{\overline{\tau}}$.

Via these results for the new construct $\widetilde{\mathbf{w}}$, we finally arrive at:

$$\mathbf{w}_{G\backslash\tau}^{HC} \ge \widetilde{\mathbf{w}}_{G\backslash\tau} = \overline{\mathbf{w}}_{G\backslash\tau} \ge \tfrac{1}{2}\mathbf{p}_{G\backslash\tau}^c + \tfrac{1}{2}\mathbf{c}_{G\backslash\tau}$$

and $\mathbf{w}_{G\cap\tau}^{HC} \ge \mathbf{p}_{G\cap\tau}^u \ge \tfrac{1}{2}\mathbf{p}_{G\cap\tau}^u + \tfrac{1}{2}\mathbf{c}_{G\cap\tau} = \tfrac{1}{2}\mathbf{p}_{G\cap\tau}^c + \tfrac{1}{2}\mathbf{c}_{G\cap\tau}$

Putting these together:

$$w_G^{HC} \geq \frac{1}{2}p_G^c + \frac{1}{2}c_G \quad \Rightarrow \quad w_G^{HC} - c_G \geq \frac{1}{2}\left(p_G^c - c_G\right)$$

■

A.0.7 Proof of Theorem 3.23

Proof. By Lemma 3.20, it only remains to consider the case when $\mathbf{p}^{1/2} \not\geq \mathbf{p}^c$. In this case, we use the newly constructed retail price $\widetilde{\mathbf{p}}^{1/2}$ and the corresponding wholesale price, $\mathbf{w}^{HC}$ to prove our proposition.

The reader may remember from discussion at the end of the lemma 3.21 that $\widetilde{\mathbf{p}}^{1/2} \geq \mathbf{p}^c$. Therefore we appeal to Lemma 3.14, which tells us that $\widetilde{p}_s^H(\mathbf{w}^{HC}) \leq p_s^H(\mathbf{w}^{HC})$. As a result,

$$\Pi^s\left(\mathbf{w}^{HC}\right) + \Pi^r\left(\mathbf{w}^{HC}\right) \geq \widetilde{\Pi}^s\left(\mathbf{w}^{HC}\right) + \Pi^r\left(\mathbf{w}^{HC}\right) = (\widetilde{\mathbf{p}}^{1/2} - \mathbf{c})^t(\mathbf{a}^c + \mathbf{u}^c - \mathbf{v}^c - \mathbf{B}\widetilde{\mathbf{p}}^{1/2}) \quad \text{(A.31)}$$

We analyze the RHS of equation (A.31) separately for the index sets G and $\overline{G}$. For $\overline{G}$, we have that:

$$\left(\widetilde{\mathbf{p}}_{\overline{G}}^{1/2} - \mathbf{c}_{\overline{G}}\right)\left[\mathbf{a}_{\overline{G}} + \mathbf{u}_{\overline{G}}^c - \mathbf{v}_{\overline{G}}^c - \mathbf{B}_{\overline{G},N}\widetilde{\mathbf{p}}^{1/2}\right]$$
$$= \left(\widetilde{\mathbf{p}}_{\overline{G}}^{1/2} - \mathbf{c}_{\overline{G}}\right)\left[\mathbf{a}_{\overline{G}}^c + \mathbf{u}_{\overline{G}}^c - \mathbf{v}_{\overline{G}}^c - \mathbf{B}_{\overline{G}}\widetilde{\mathbf{p}}_{\overline{G}}^{1/2}\right]$$

To establish a lower bound on the above expression, we argue the applicability of Lemma 3.18 to it. Suppose we were to assume that all retailers in the index set G *fixed* their retail prices at $\mathbf{p}_G^c$ apriori. In this case, the optimal centralized retail price would be $\mathbf{p}^G$ (by definition) and expression (3.8) would reduce to expression (3.13). Hence it follows from Lemma 3.18 and from a similar exercise as in Lemma 3.20 that:

$$\geq \frac{3}{4}\left(\mathbf{p}_{\overline{G}}^G - \mathbf{c}_{\overline{G}}\right)\left[\mathbf{a}_{\overline{G}}^c + \mathbf{u}_{\overline{G}}^c - \mathbf{v}_{\overline{G}}^c - \mathbf{B}_{\overline{G}}\mathbf{p}_{\overline{G}}^G\right] = \frac{3}{4}\left(\mathbf{p}_{\overline{G}}^G - \mathbf{c}_{\overline{G}}\right)\left[\mathbf{a}_{\overline{G}} + \mathbf{u}_{\overline{G}}^c - \mathbf{v}_{\overline{G}}^c - \mathbf{B}_{\overline{G},N}\mathbf{p}^G\right]$$

$$\text{(A.32)}$$

For the index set G, we know that $\widetilde{\mathbf{p}}_G^{1/2} = \mathbf{p}_G^G = \mathbf{p}_G^c$ and so we have that:

$$\left(\widetilde{\mathbf{p}}_G^{1/2} - \mathbf{c}_G\right)\left[\mathbf{a}_G + \mathbf{u}_G^c - \mathbf{v}_G^c - \mathbf{B}_{G,N}\widetilde{\mathbf{p}}^{1/2}\right] = \left(\mathbf{p}_G^G - \mathbf{c}_G\right)\left[\mathbf{a}_G + \mathbf{u}_G^c - \mathbf{v}_G^c - \mathbf{B}_{G,N}\widetilde{\mathbf{p}}^{1/2}\right]$$

$$\geq \left(\mathbf{p}_G^G - \mathbf{c}_G\right)\left[\mathbf{a}_G + \mathbf{u}_G^c - \mathbf{v}_G^c - \mathbf{B}_{G,N}\mathbf{p}^G\right].$$

Putting together the results for G and $\overline{G}$ in expression (A.31):

$$\Pi^s\left(\mathbf{w}^{HC}\right) + \Pi^r\left(\mathbf{w}^{HC}\right) \geq \frac{3}{4}\left(\mathbf{p}^G - \mathbf{c}\right)\left[\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c - \mathbf{B}\mathbf{p}^G\right] = \frac{3}{4}\Pi^r\left(\mathbf{c}, \mathbf{p}^G\right) \geq \frac{3}{4}\Pi^c\left(\mathbf{p}^c, \mathbf{u}^c\right)$$

Observe that the last inequality is a consequence of Lemma 3.13.

For the supplier's profit, we have that:

$$\Pi^s\left(\mathbf{w}^{HC}\right) \geq \widetilde{\Pi}^s\left(\mathbf{w}^{HC}\right) = \left(\mathbf{w}^{HC} - \mathbf{c}\right)^t\left[\mathbf{a} - \mathbf{B}\widetilde{\mathbf{p}}^{1/2} + \mathbf{u}^c - \mathbf{v}^c\right]$$

To compare this expression with $p_c(\mathbf{p}^c, \mathbf{u}^c)$, we need to replace the terms $\widetilde{\mathbf{p}}^{1/2}$ and $\mathbf{w}^{HC}$ appropriately. So, we split the expression into three parts, corresponding to the index sets G, $\varsigma(1/2)$ and $\widetilde{\varsigma(1/2)}$. It follows then follows from $\varsigma(1/2) \subseteq \tau$ that:

$$= \left(\mathbf{w}_G^{HC} - \mathbf{c}_G\right)^t\left(\mathbf{a}_G + \mathbf{u}_G^c - \mathbf{v}_G^c - \mathbf{B}_{G,N}\widetilde{\mathbf{p}}^{1/2}\right)$$

$$+ \left(\mathbf{p}_{\varsigma(1/2)}^u - \mathbf{c}_{\varsigma(1/2)}\right)^t\left(\mathbf{a}_{\varsigma(1/2)} + \mathbf{u}_{\varsigma(1/2)}^c - \mathbf{v}_{\varsigma(1/2)}^c - \mathbf{B}_{\varsigma(1/2),N}\widetilde{\mathbf{p}}^{1/2}\right)$$

$$+ \left(\mathbf{w}_{\widetilde{\varsigma(1/2)}}^{HC} - \mathbf{c}_{\widetilde{\varsigma(1/2)}}\right)^t\left(\mathbf{a}_{\widetilde{\varsigma(1/2)}} + \mathbf{u}_{\widetilde{\varsigma(1/2)}}^c - \mathbf{v}_{\widetilde{\varsigma(1/2)}}^c - \mathbf{B}_{\widetilde{\varsigma(1/2)},N}\widetilde{\mathbf{p}}^{1/2}\right)$$

$$= \left(\mathbf{w}_G^{HC} - \mathbf{c}_G\right)^t\left(\mathbf{a}_G + \mathbf{u}_G^c - \mathbf{v}_G^c - \mathbf{B}_{G,N}\mathbf{p}^G\right) - \left(\mathbf{w}_G^{HC} - \mathbf{c}_G\right)^t\mathbf{B}_{G,\overline{G}}\left(\widetilde{\mathbf{p}}_{\overline{G}}^{1/2} - \mathbf{p}_{\overline{G}}^G\right)$$

$$+ \left(\mathbf{p}_{\varsigma(1/2)}^u - \mathbf{c}_{\varsigma(1/2)}\right)^t\left(\mathbf{a}_{\varsigma(1/2)}^c + \mathbf{u}_{\varsigma(1/2)}^c - \mathbf{v}_{\varsigma(1/2)}^c - \mathbf{B}_{\varsigma(1/2),\overline{G}}\widetilde{\mathbf{p}}_{\overline{G}}^{1/2}\right)$$

$$+ \left(\mathbf{w}_{\widetilde{\varsigma(1/2)}}^{HC} - \mathbf{c}_{\widetilde{\varsigma(1/2)}}\right)^t\left(\mathbf{a}_{\widetilde{\varsigma(1/2)}}^c + \mathbf{u}_{\widetilde{\varsigma(1/2)}}^c - \mathbf{v}_{\widetilde{\varsigma(1/2)}}^c - \mathbf{B}_{\widetilde{\varsigma(1/2)},\overline{G}}\widetilde{\mathbf{p}}_{\overline{G}}^{1/2}\right)$$

Further, from Lemma 3.22, and the optimality conditions of expressions (3.14) and (3.15):

$$\geq \frac{1}{2}\left(\mathbf{p}_G^c - \mathbf{c}_G\right)^t\left(\mathbf{a}_G + \mathbf{u}_G^c - \mathbf{v}_G^c - \mathbf{B}_{G,N}\mathbf{p}^G\right) - \left(\mathbf{w}_G^{HC} - \mathbf{c}_G\right)^t\mathbf{B}_{G,\overline{G}}\left(\widetilde{\mathbf{p}}_{\overline{G}}^{1/2} - \mathbf{p}_{\overline{G}}^G\right)$$

$$+ \frac{1}{2}\left(\mathbf{p}_{\varsigma(1/2)}^u - \mathbf{c}_{\varsigma(1/2)}\right)^t\left(\mathbf{a}_{\varsigma(1/2)}^c + \mathbf{u}_{\varsigma(1/2)}^c - \mathbf{v}_{\varsigma(1/2)}^c - \mathbf{B}_{\varsigma(1/2),\overline{G}}\mathbf{p}_{\overline{G}}^G\right)$$

$$+ \frac{1}{2}\left(\mathbf{w}_{\widetilde{\varsigma(1/2)}}^{HC} - \mathbf{c}_{\widetilde{\varsigma(1/2)}}\right)^t\left(\mathbf{a}_{\widetilde{\varsigma(1/2)}}^c + \mathbf{u}_{\widetilde{\varsigma(1/2)}}^c - \mathbf{v}_{\widetilde{\varsigma(1/2)}}^c - \mathbf{B}_{\widetilde{\varsigma(1/2)},\overline{G}}\mathbf{p}_{\overline{G}}^G\right) \qquad \text{(A.33)}$$

We now derive a lower bound on $\mathbf{w}_{\widetilde{\varsigma(1/2)}}^{HC}$ to be substituted in inequality (A.33). Rearranging the definition of $\mathbf{w}_{\widetilde{\tau}}^{HC}$ (in a similar manner as Lemma 3.19) and restricting to the subset $\widetilde{\varsigma(1/2)} \subseteq \widetilde{\tau}$, we have that:

$$0 = \mathbf{a}_{\widetilde{\varsigma(1/2)}} + \mathbf{u}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{v}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{B}_{\widetilde{\varsigma(1/2)},N}\widetilde{\mathbf{p}}^{1/2} - \mathbf{B}_{\widetilde{\varsigma(1/2)},\overline{T}}\left(\widetilde{\mathbf{p}}^{1/2}_{\overline{T}} - \mathbf{w}^{HC}_{\overline{T}}\right)$$

$$= \mathbf{a}^c_{\widetilde{\varsigma(1/2)}} + \mathbf{u}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{v}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{B}_{\widetilde{\varsigma(1/2)},\overline{G}}\widetilde{\mathbf{p}}^{1/2}_G - \mathbf{B}_{\widetilde{\varsigma(1/2)}}\left(\widetilde{\mathbf{p}}^{1/2}_{\widetilde{\varsigma(1/2)}} - \mathbf{w}^{HC}_{\widetilde{\varsigma(1/2)}}\right)$$

$$- \mathbf{B}_{\widetilde{\varsigma(1/2)},\overline{T}\backslash\widetilde{\varsigma(1/2)}}\left(\widetilde{\mathbf{p}}^{1/2}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}} - \mathbf{w}^{HC}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}}\right)$$

Rearranging this for $\mathbf{w}^{HC}_{\widetilde{\varsigma(1/2)}}$:

$$\mathbf{w}^{HC}_{\widetilde{\varsigma(1/2)}} = 2\widetilde{\mathbf{p}}^{1/2}_{\widetilde{\varsigma(1/2)}} - \mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\left(\mathbf{a}^c_{\widetilde{\varsigma(1/2)}} + \mathbf{u}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{v}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{B}_{\widetilde{\varsigma(1/2)},\varsigma(1/2)}\mathbf{p}^u_{\varsigma(1/2)}\right)$$

$$+ \mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\mathbf{B}_{\widetilde{\varsigma(1/2)},\overline{T}\backslash\widetilde{\varsigma(1/2)}}\left(\widetilde{\mathbf{p}}^{1/2}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}} - \mathbf{w}^{HC}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}}\right)$$

From the optimality condition (3.14), we have that:

$$\mathbf{a}^c_{\widetilde{\varsigma(1/2)}} + \mathbf{u}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{v}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{B}_{\widetilde{\varsigma(1/2)},\overline{G}}\widetilde{\mathbf{p}}^{1/2}_G = \tfrac{1}{2}\left(\mathbf{a}^c_{\widetilde{\varsigma(1/2)}} + \mathbf{u}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{v}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{B}_{\widetilde{\varsigma(1/2)},\overline{G}}\mathbf{p}^G_G\right)$$

Since $\widetilde{\varsigma(1/2)} \subseteq \overline{G}$, we may rearrange the optimality condition above to derive an expression for $\widetilde{\mathbf{p}}^{1/2}_{\widetilde{\varsigma(1/2)}}$. Substituting this expression in the above formula for $\mathbf{w}^{HC}_{\widetilde{\varsigma(1/2)}}$, we then have that:

$$\mathbf{w}^{HC}_{\widetilde{\varsigma(1/2)}} = \mathbf{p}^G_{\widetilde{\varsigma(1/2)}} + \mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\left(\mathbf{a}^c_{\widetilde{\varsigma(1/2)}} + \mathbf{u}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{v}^c_{\widetilde{\varsigma(1/2)}}\right) + \mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\mathbf{B}_{\widetilde{\varsigma(1/2)},\varsigma(1/2)}\mathbf{p}^G_{\varsigma(1/2)}$$

$$- 2\mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\mathbf{B}_{\widetilde{\varsigma(1/2)},\varsigma(1/2)}\mathbf{p}^u_{\varsigma(1/2)}$$

$$- \mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\left(\mathbf{a}^c_{\widetilde{\varsigma(1/2)}} + \mathbf{u}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{v}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{B}_{\widetilde{\varsigma(1/2)},\varsigma(1/2)}\mathbf{p}^u_{\varsigma(1/2)}\right)$$

$$+ \mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\mathbf{B}_{\widetilde{\varsigma(1/2)},\overline{T}\backslash\widetilde{\varsigma(1/2)}}\left(\widetilde{\mathbf{p}}^{1/2}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}} - \mathbf{w}^{HC}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}}\right)$$

$$= \mathbf{p}^G_{\widetilde{\varsigma(1/2)}} - \mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\mathbf{B}_{\widetilde{\varsigma(1/2)},\varsigma(1/2)}\left(\mathbf{p}^u_{\varsigma(1/2)} - \mathbf{p}^G_{\varsigma(1/2)}\right)$$

$$+ \mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\mathbf{B}_{\widetilde{\varsigma(1/2)},\overline{T}\backslash\widetilde{\varsigma(1/2)}}\left(\widetilde{\mathbf{p}}^{1/2}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}} - \mathbf{w}^{HC}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}}\right)$$

$$\geq \mathbf{p}^G_{\widetilde{\varsigma(1/2)}} + \mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\mathbf{B}_{\widetilde{\varsigma(1/2)},\overline{T}\backslash\widetilde{\varsigma(1/2)}}\left(\widetilde{\mathbf{p}}^{1/2}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}} - \mathbf{w}^{HC}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}}\right)$$

Substituting back this lower bound on $\mathbf{w}^{HC}_{\widetilde{\varsigma(1/2)}}$ in expression (A.33), we have that:

$$\geq \tfrac{1}{2}\left(\mathbf{p}^c_G - \mathbf{c}_G\right)^t\left(\mathbf{a}_G + \mathbf{u}^c_G - \mathbf{v}^c_G - \mathbf{B}_{G,N}\mathbf{p}^G\right) - \left(\mathbf{w}^{HC}_G - \mathbf{c}_G\right)^t\mathbf{B}_{G,\overline{G}}\left(\widetilde{\mathbf{p}}^{1/2}_{\overline{G}} - \mathbf{p}^G_{\overline{G}}\right)$$

$$+ \tfrac{1}{2}\left(\mathbf{p}^u_{\varsigma(1/2)} - \mathbf{c}_{\varsigma(1/2)}\right)^t\left(\mathbf{a}^c_{\varsigma(1/2)} + \mathbf{u}^c_{\varsigma(1/2)} - \mathbf{v}^c_{\varsigma(1/2)} - \mathbf{B}_{\varsigma(1/2),\overline{G}}\mathbf{p}^G_{\overline{G}}\right)$$

$$+\tfrac{1}{2}\left(\mathbf{p}^G_{\widetilde{\varsigma(1/2)}} - \mathbf{c}_{\widetilde{\varsigma(1/2)}}\right)^t \left(\mathbf{a}^c_{\widetilde{\varsigma(1/2)}} + \mathbf{u}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{v}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{B}_{\widetilde{\varsigma(1/2)},\overline{G}}\,\mathbf{p}^G_G\right)$$

$$+\frac{1}{2}\left(\widetilde{\mathbf{p}}^{1/2}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}} - \mathbf{w}^{HC}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}}\right)^t \mathbf{B}_{\overline{T}\backslash\widetilde{\varsigma(1/2)},\widetilde{\varsigma(1/2)}} \mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\left(\mathbf{a}^c_{\widetilde{\varsigma(1/2)}} + \mathbf{u}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{v}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{B}_{\widetilde{\varsigma(1/2)},\overline{G}}\mathbf{p}^G_G\right)$$

$$\tag{A.34}$$

Since $\widetilde{\mathbf{p}}^{1/2}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}} - \mathbf{w}^{HC}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}} \leq \mathbf{w}^{HC}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}} - \mathbf{c}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}}$ by lemma 3.22, we have for the last part of expression (A.34) that:

$$\left(\widetilde{\mathbf{p}}^{1/2}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}} - \mathbf{w}^{HC}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}}\right)^t \mathbf{B}_{\overline{T}\backslash\widetilde{\varsigma(1/2)},\widetilde{\varsigma(1/2)}} \mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\left(\mathbf{a}^c_{\widetilde{\varsigma(1/2)}} + \mathbf{u}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{v}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{B}_{\widetilde{\varsigma(1/2)},\overline{G}}\mathbf{p}^G_G\right)$$

$$\geq \left(\mathbf{w}^{HC}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}} - \mathbf{c}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}}\right)^t \mathbf{B}_{\overline{T}\backslash\widetilde{\varsigma(1/2)},\widetilde{\varsigma(1/2)}} \mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\left(\mathbf{a}^c_{\widetilde{\varsigma(1/2)}} + \mathbf{u}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{v}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{B}_{\widetilde{\varsigma(1/2)},\overline{G}}\mathbf{p}^G_G\right)$$

$$= \left(\mathbf{w}^{HC}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}} - \mathbf{c}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}}\right)^t \mathbf{B}_{\overline{T}\backslash\widetilde{\varsigma(1/2)},\widetilde{\varsigma(1/2)}} \mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\mathbf{B}_{\widetilde{\varsigma(1/2)},\overline{G}}\left(\mathbf{p}^G_G - \mathbf{c}_{\overline{G}}\right)$$

$$= \left(\mathbf{w}^{HC}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}} - \mathbf{c}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}}\right)^t \mathbf{B}_{\overline{T}\backslash\widetilde{\varsigma(1/2)},\widetilde{\varsigma(1/2)}}\left(\mathbf{p}^G_{\widetilde{\varsigma(1/2)}} - \mathbf{c}_{\widetilde{\varsigma(1/2)}}\right)$$

$$+ \left(\mathbf{w}^{HC}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}} - \mathbf{c}_{\overline{T}\backslash\widetilde{\varsigma(1/2)}}\right)^t \mathbf{B}_{\overline{T}\backslash\widetilde{\varsigma(1/2)},\widetilde{\varsigma(1/2)}} \mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\mathbf{B}_{\widetilde{\varsigma(1/2)},\varsigma(1/2)}\left(\mathbf{p}^G_{\varsigma(1/2)} - \mathbf{c}_{\varsigma(1/2)}\right) \tag{A.35}$$

Moreover, the optimality condition of equation (3.14) implies that:

$$\widetilde{\mathbf{p}}^{1/2}_{\widetilde{\varsigma(1/2)}} - \mathbf{p}^G_{\widetilde{\varsigma(1/2)}}$$

$$= \tfrac{1}{2}\mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\left(\mathbf{a}^c_{\widetilde{\varsigma(1/2)}} + \mathbf{u}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{v}^c_{\widetilde{\varsigma(1/2)}} - \mathbf{B}_{\widetilde{\varsigma(1/2)},\overline{G}}\mathbf{p}^G_G\right) - \mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\mathbf{B}_{\widetilde{\varsigma(1/2)},\varsigma(1/2)}\left(\widetilde{\mathbf{p}}^{1/2}_{\varsigma(1/2)} - \mathbf{p}^G_{\varsigma(1/2)}\right)$$

Using the optimality condition of $\mathbf{p}^G$ for the index set $\widetilde{\varsigma(1/2)}$:

$$= \tfrac{1}{2}\mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\mathbf{B}_{\widetilde{\varsigma(1/2)},\overline{G}}\left(\mathbf{p}^G_G - \mathbf{c}_{\overline{G}}\right) - \mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\mathbf{B}_{\widetilde{\varsigma(1/2)},\varsigma(1/2)}\left(\widetilde{\mathbf{p}}^{1/2}_{\varsigma(1/2)} - \mathbf{p}^G_{\varsigma(1/2)}\right)$$

$$= \tfrac{1}{2}\left(\mathbf{p}^G_{\widetilde{\varsigma(1/2)}} - \mathbf{c}_{\widetilde{\varsigma(1/2)}}\right) + \tfrac{1}{2}\mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\mathbf{B}_{\widetilde{\varsigma(1/2)},\varsigma(1/2)}\left(\mathbf{p}^G_{\varsigma(1/2)} - \mathbf{c}_{\varsigma(1/2)}\right)$$

$$\quad - \mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\mathbf{B}_{\widetilde{\varsigma(1/2)},\varsigma(1/2)}\left(\widetilde{\mathbf{p}}^{1/2}_{\varsigma(1/2)} - \mathbf{p}^G_{\varsigma(1/2)}\right)$$

$$\geq \tfrac{1}{2}\left(\mathbf{p}^G_{\widetilde{\varsigma(1/2)}} - \mathbf{c}_{\widetilde{\varsigma(1/2)}}\right) + \tfrac{1}{2}\mathbf{B}^{-1}_{\widetilde{\varsigma(1/2)}}\mathbf{B}_{\widetilde{\varsigma(1/2)},\varsigma(1/2)}\left(\mathbf{p}^G_{\varsigma(1/2)} - \mathbf{c}_{\varsigma(1/2)}\right).$$

Since by lemma 3.22 $\mathbf{w}^{HC}_G \geq \mathbf{c}_G$, and by lemma 3.21 $\widetilde{\mathbf{p}}^{1/2}_{\overline{G}} \geq \mathbf{p}^G_{\overline{G}}$, we may use the previous result to bound the second part of expression (A.34) as:

$$- \left(\mathbf{w}^{HC}_G - \mathbf{c}_G\right)^t \mathbf{B}_{G,\overline{G}}\left(\widetilde{\mathbf{p}}^{1/2}_{\overline{G}} - \mathbf{p}^G_{\overline{G}}\right)$$

$$\geq -\left(\mathbf{w}^{HC}_{\overline{\mathcal{T}}\backslash\varsigma(1/2)} - \mathbf{c}_{\overline{\mathcal{T}}\backslash\varsigma(1/2)}\right)^t \mathbf{B}_{\overline{\mathcal{T}}\backslash\varsigma(1/2),\varsigma(1/2)} \left(\widetilde{\mathbf{p}}^{1/2}_{\varsigma(1/2)} - \mathbf{p}^G_{\varsigma(1/2)}\right)$$

$$\geq -\tfrac{1}{2}\left(\mathbf{w}^{HC}_{\overline{\mathcal{T}}\backslash\varsigma(1/2)} - \mathbf{c}_{\overline{\mathcal{T}}\backslash\varsigma(1/2)}\right)^t \mathbf{B}_{\overline{\mathcal{T}}\backslash\varsigma(1/2),\varsigma(1/2)} \left(\mathbf{p}^G_{\varsigma(1/2)} - \mathbf{c}_{\varsigma(1/2)}\right)$$

$$-\tfrac{1}{2}\left(\mathbf{w}^{HC}_{\overline{\mathcal{T}}\backslash\varsigma(1/2)} - \mathbf{c}_{\overline{\mathcal{T}}\backslash\varsigma(1/2)}\right)^t \mathbf{B}_{\overline{\mathcal{T}}\backslash\varsigma(1/2),\varsigma(1/2)}\mathbf{B}^{-1}_{\varsigma(1/2)}\mathbf{B}_{\varsigma(1/2),\varsigma(1/2)} \left(\mathbf{p}^G_{\varsigma(1/2)} - \mathbf{c}_{\varsigma(1/2)}\right)$$

Notice that this expression is exactly the RHS of inequality (A.35) and so it follows that:

$$\geq -\frac{1}{2}\left(\widetilde{\mathbf{p}}^{1/2}_{\overline{\mathcal{T}}\backslash\varsigma(1/2)} - \mathbf{w}^{HC}_{\overline{\mathcal{T}}\backslash\varsigma(1/2)}\right)\mathbf{B}_{\overline{\mathcal{T}}\backslash\varsigma(1/2),\varsigma(1/2)}\mathbf{B}^{-1}_{\varsigma(1/2)}\left(\mathbf{a}^c_{\varsigma(1/2)} + \mathbf{u}^c_{\varsigma(1/2)} - \mathbf{v}^c_{\varsigma(1/2)} - \mathbf{B}_{\varsigma(1/2),\overline{G}}\mathbf{p}^G_{\overline{G}}\right)$$

$$\tag{A.36}$$

Finally, substituting inequality (A.36) in expression (A.34):

$$\Pi^s\left(\mathbf{w}^{HC}\right) \geq \tfrac{1}{2}\left(\mathbf{p}^c_G - \mathbf{c}_G\right)^t\left(\mathbf{a}_G + \mathbf{u}^c_G - \mathbf{v}^c_G - \mathbf{B}_{G,N}\mathbf{p}^G\right)$$

$$+\tfrac{1}{2}\left(\mathbf{p}^u_{\varsigma(1/2)} - \mathbf{c}_{\varsigma(1/2)}\right)^t\left(\mathbf{a}^c_{\varsigma(1/2)} + \mathbf{u}^c_{\varsigma(1/2)} - \mathbf{v}^c_{\varsigma(1/2)} - \mathbf{B}_{\varsigma(1/2),\overline{G}}\mathbf{p}^G_{\overline{G}}\right)$$

$$+\tfrac{1}{2}\left(\mathbf{p}^G_{\varsigma(1/2)} - \mathbf{c}_{\varsigma(1/2)}\right)^t\left(\mathbf{a}^c_{\varsigma(1/2)} + \mathbf{u}^c_{\varsigma(1/2)} - \mathbf{v}^c_{\varsigma(1/2)} - \mathbf{B}_{\varsigma(1/2),\overline{G}}\mathbf{p}^G_{\overline{G}}\right)$$

$$\geq \tfrac{1}{2}\left(\mathbf{p}^G_G - \mathbf{c}_G\right)^t\left(\mathbf{a}_G + \mathbf{u}^c_G - \mathbf{v}^c_G - \mathbf{B}_{G,N}\mathbf{p}^G\right)$$

$$+\tfrac{1}{2}\left(\mathbf{p}^G_{\varsigma(1/2)} - \mathbf{c}_{\varsigma(1/2)}\right)^t\left(\mathbf{a}^c_{\varsigma(1/2)} + \mathbf{u}^c_{\varsigma(1/2)} - \mathbf{v}^c_{\varsigma(1/2)} - \mathbf{B}_{\varsigma(1/2),\overline{G}}\mathbf{p}^G_{\overline{G}}\right)$$

$$+\tfrac{1}{2}\left(\mathbf{p}^G_{\varsigma(1/2)} - \mathbf{c}_{\varsigma(1/2)}\right)^t\left(\mathbf{a}^c_{\varsigma(1/2)} + \mathbf{u}^c_{\varsigma(1/2)} - \mathbf{v}^c_{\varsigma(1/2)} - \mathbf{B}_{\varsigma(1/2),\overline{G}}\mathbf{p}^G_{\overline{G}}\right)$$

$$= \tfrac{1}{2}\left(\mathbf{p}^G - \mathbf{c}\right)^t\left(\mathbf{a} + \mathbf{u}^c - \mathbf{v}^c - \mathbf{B}\mathbf{p}^G\right) = \tfrac{1}{2}\Pi^r\left(\mathbf{c},\mathbf{p}^G\right) \geq \tfrac{1}{2}\Pi^c\left(\mathbf{p}^c,\mathbf{u}^c\right). \qquad \blacksquare$$

A.0.8 Proof of Lemma 3.25

Proof. Suppose that the boundary of $\mathbf{p}(\widehat{\mathbf{w}})$ is indeed $L(\mathbf{w})$, i.e., $L^e(\widehat{\mathbf{w}}) = L(\mathbf{w})$. We then show using our definition of $\widehat{\mathbf{w}}$ that all optimality conditions are satisfied by $\mathbf{p}(\widehat{\mathbf{w}})$ at this boundary.

First note that $\mathbf{p}_{L(\mathbf{w})}(\widehat{\mathbf{w}}) = \mathbf{p}^u_{L(\mathbf{w})}$, by our assumption. Then from the optimality condition of equation (3.18), it follows that:

$$\mathbf{p}_{\overline{L(\mathbf{w})}}(\widehat{\mathbf{w}}) = \left(\mathbf{B}_{\overline{L(\mathbf{w})}} + \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}}\right)^{-1}\left(\mathbf{a}_{\overline{L(\mathbf{w})}} + \mathbf{u}^c_{\overline{L(\mathbf{w})}} - \mathbf{v}^c_{\overline{L(\mathbf{w})}} + \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})},N}\widehat{\mathbf{w}}\right)$$

$$- \left(\mathbf{B}_{\overline{L(\mathbf{w})}} + \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}}\right)^{-1}\left(\mathbf{B}_{\overline{L(\mathbf{w})},L(\mathbf{w})} + \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})},L(\mathbf{w})}\right)\mathbf{p}^u_{L(\mathbf{w})}$$

$$= \left(\mathbf{B}_{\overline{L(\mathbf{w})}} + \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}}\right)^{-1}\left(\mathbf{a}_{\overline{L(\mathbf{w})}} + \mathbf{u}^c_{\overline{L(\mathbf{w})}} - \mathbf{v}^c_{\overline{L(\mathbf{w})}} + \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}}\widehat{\mathbf{w}}_{\overline{L(\mathbf{w})}}\right)$$

$$- \left(\mathbf{B}_{\overline{L(\mathbf{w})}} + \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}}\right)^{-1}\mathbf{B}_{\overline{L(\mathbf{w})},L(\mathbf{w})}\mathbf{p}^u_{L(\mathbf{w})} \qquad\qquad \left(\because \widehat{\mathbf{w}}_{L(\mathbf{w})} = \mathbf{p}^u_{L(\mathbf{w})}\right)$$

$$= \left(\mathbf{B}_{\overline{L(\mathbf{w})}} + \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}}\right)^{-1} \mathbf{B}_{\overline{L(\mathbf{w})}} \mathbf{B}_{\overline{L(\mathbf{w})}}^{-1} \left(\mathbf{a}_{\overline{L(\mathbf{w})}} + \mathbf{u}_{\overline{L(\mathbf{w})}}^c - \mathbf{v}_{\overline{L(\mathbf{w})}}^c\right)$$

$$+ \left(\mathbf{B}_{\overline{L(\mathbf{w})}} + \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}}\right)^{-1} \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}} \left[\mathbf{w}_{\overline{L(\mathbf{w})}} - \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}}^{-1} \left(\mathbf{B}_{\overline{L(\mathbf{w})}} - \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}}\right) \left(\mathbf{p}_{\overline{L(\mathbf{w})}}(\mathbf{w}) - \mathbf{w}_{\overline{L(\mathbf{w})}}\right)\right]$$

$$- \left(\mathbf{B}_{\overline{L(\mathbf{w})}} + \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}}\right)^{-1} \mathbf{B}_{\overline{L(\mathbf{w})}, L(\mathbf{w})} \left(\mathbf{p}_{L(\mathbf{w})}^u - \mathbf{w}_{L(\mathbf{w})}\right)$$

$$- \left(\mathbf{B}_{\overline{L(\mathbf{w})}} + \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}}\right)^{-1} \mathbf{B}_{\overline{L(\mathbf{w})}, L(\mathbf{w})} \mathbf{p}_{L(\mathbf{w})}^u \qquad \text{(substituting for} \mathbf{p}_{\overline{L(\mathbf{w})}}(\widehat{\mathbf{w}}))$$

Similarly, substituting for $\mathbf{B}_{\overline{L(\mathbf{w})}}^{-1} \left(\mathbf{a}_{\overline{L(\mathbf{w})}} + \mathbf{u}_{\overline{L(\mathbf{w})}}^c - \mathbf{v}_{\overline{L(\mathbf{w})}}^c\right)$ from the optimality condition of $\mathbf{p}(\mathbf{w})$ for the index set $\overline{L(\mathbf{w})}$ (see for e.c., equation (3.5)):

$$= \left(\mathbf{B}_{\overline{L(\mathbf{w})}} + \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}}\right)^{-1} \mathbf{B}_{\overline{L(\mathbf{w})}} \left[\left(2\mathbf{p}_{\overline{L(\mathbf{w})}}(\mathbf{w}) - \mathbf{w}_{\overline{L(\mathbf{w})}}\right) + \mathbf{B}_{\overline{L(\mathbf{w})}}^{-1} \mathbf{B}_{\overline{L(\mathbf{w})}, L(\mathbf{w})} \left(2\mathbf{p}_{L(\mathbf{w})}^u - \mathbf{w}_{L(\mathbf{w})}\right)\right]$$

$$+ \left(\mathbf{B}_{\overline{L(\mathbf{w})}} + \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}}\right)^{-1} \left[\widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}} \mathbf{w}_{\overline{L(\mathbf{w})}} - \left(\mathbf{B}_{\overline{L(\mathbf{w})}} - \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}}\right) \left(\mathbf{p}_{\overline{L(\mathbf{w})}}(\mathbf{w}) - \mathbf{w}_{\overline{L(\mathbf{w})}}\right)\right]$$

$$- \left(\mathbf{B}_{\overline{L(\mathbf{w})}} + \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}}\right)^{-1} \mathbf{B}_{\overline{L(\mathbf{w})}, L(\mathbf{w})} \left(2\mathbf{p}_{L(\mathbf{w})}^u - \mathbf{w}_{L(\mathbf{w})}\right)$$

$$= \left(\mathbf{B}_{\overline{L(\mathbf{w})}} + \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}}\right)^{-1} \left(\mathbf{B}_{\overline{L(\mathbf{w})}} + \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}}\right) \mathbf{p}_{\overline{L(\mathbf{w})}}(\mathbf{w}) = \mathbf{p}_{\overline{L(\mathbf{w})}}(\mathbf{w}) < \mathbf{p}_{\overline{L(\mathbf{w})}}^u$$

Hence we have that $\mathbf{p}_{\overline{L(\mathbf{w})}}(\widehat{\mathbf{w}}) = \mathbf{p}_{\overline{L(\mathbf{w})}}(\mathbf{w}) < \mathbf{p}_{\overline{L(\mathbf{w})}}^u$ and so all optimality conditions are satisfied for the index set $\overline{L(\mathbf{w})}$. To verify the optimality condition of equation (3.17), we see that:

$$\mathbf{a}_{L(\mathbf{w})} + \mathbf{u}_{L(\mathbf{w})}^c - \mathbf{v}_{L(\mathbf{w})}^c - \mathbf{B}_{L(\mathbf{w}), N} \mathbf{p}(\mathbf{w}) - \widehat{\mathbf{B}}_{L(\mathbf{w}), N} \left[\mathbf{p}(\mathbf{w}) - \widehat{\mathbf{w}}\right]$$

$$= \mathbf{a}_{L(\mathbf{w})} + \mathbf{u}_{L(\mathbf{w})}^c - \mathbf{v}_{L(\mathbf{w})}^c - \mathbf{B}_{L(\mathbf{w}), N} \mathbf{p}(\mathbf{w}) \qquad (\because (\widehat{\mathbf{w}})_{L(\mathbf{w})} = \mathbf{p}_{L(\mathbf{w})}^u = \mathbf{p}_{L(\mathbf{w})}(\mathbf{w}))$$

$$\geq \mathbf{a}_{L(\mathbf{w})} - \mathbf{B}_{L(\mathbf{w}), N} \mathbf{p}(\mathbf{w}) \geq 0.$$

Thus we have verified that all optimality conditions are satisfied at $L(\mathbf{w})$ and moreover, that $\mathbf{p}^e(\widehat{\mathbf{w}}) = \mathbf{p}(\mathbf{w})$.

Note that by the optimality condition of equation (3.18):

$$0 = \mathbf{a}_{\overline{L(\mathbf{w})}} + \mathbf{u}_{\overline{L(\mathbf{w})}}^c - \mathbf{v}_{\overline{L(\mathbf{w})}}^c - \mathbf{B}_{\overline{L(\mathbf{w})}, N} \mathbf{p}(\mathbf{w}) - \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}, N} \left[\mathbf{p}(\mathbf{w}) - \widehat{\mathbf{w}}\right]$$

$$= \mathbf{a}_{\overline{L(\mathbf{w})}} + \mathbf{u}_{\overline{L(\mathbf{w})}}^c - \mathbf{v}_{\overline{L(\mathbf{w})}}^c - \mathbf{B}_{\overline{L(\mathbf{w})}, N} \mathbf{p}(\mathbf{w}) - \widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}} \left[\mathbf{p}_{\overline{L(\mathbf{w})}}(\mathbf{w}) - \widehat{\mathbf{w}}_{\overline{L(\mathbf{w})}}\right]$$

and because $\mathbf{a}_{\overline{L(\mathbf{w})}} + \mathbf{u}_{\overline{L(\mathbf{w})}}^c - \mathbf{v}_{\overline{L(\mathbf{w})}}^c - \mathbf{B}_{\overline{L(\mathbf{w})}, N} \mathbf{p}(\mathbf{w}) \geq \mathbf{a}_{\overline{L(\mathbf{w})}} - \mathbf{B}_{\overline{L(\mathbf{w})}, N} \mathbf{p}(\mathbf{w}) \geq 0$

it must be that $\widehat{\mathbf{B}}_{\overline{L(\mathbf{w})}} \left[\mathbf{p}_{\overline{L(\mathbf{w})}}(\mathbf{w}) - \widehat{\mathbf{w}}_{\overline{L(\mathbf{w})}}\right] \geq 0$, and hence $\mathbf{p}_{\overline{L(\mathbf{w})}}(\mathbf{w}) \geq \widehat{\mathbf{w}}_{\overline{L(\mathbf{w})}}$. ∎

A.0.9 Proof of Theorem 3.30

Proof. In a coordinated system, suppose the optimal retail price vector is $\mathbf{p}^o$ and the corresponding optimal inventory vector is $\mathbf{a} - \mathbf{B}\mathbf{p}^o + \mathbf{u}^o$ for the wholesale price vector $\mathbf{w}^o$ and

buy-back menu $s^o(w^o, \bullet)$. For $i = 1, 2, ..., n$, the optimality of the inventory levels implies that $\frac{p_i^o - w_i^o}{p_i^o - s_i^o(w_i^o, p_i^o)} = F_i(u_i^o)$. Let $o_i^o = E[u_i^o - \varepsilon_i]^+$ and $v_i^o = \frac{E[u_i^o - \varepsilon_i]^+}{F_i(u_i^o)}$. Now, the expected retailer profit may be expressed as:

$$\Pi^r(w^o, s^c(w^o, \bullet), p^o, u^o) = (p^o - w^o)^t(a - Bp^o) + (p^o - w^o)^t u^o - [p^o - s^o(w^o, p^o)]^t o^o$$

$$= (p^o - w^o)^t(a - Bp^o + u^o - v^o).$$

In addition, suppose that $L^o = \{i \in N : p_i^o = p_i^u\}$. The optimality of p^o for the retailer would imply that:

$$a_{L^o} - B_{L^o, N}p^o + u_{L^o}^o - v_{L^o}^o - B_{L^o, N}(p^o - w^o) \geq 0$$

and

$$a_{\overline{L^o}} - B_{\overline{L^o}, N}p^o + u_{\overline{L^o}}^o - v_{\overline{L^o}}^o - B_{\overline{L^o}, N}(p^o - w^o) = 0.$$

In a decentralized setting, consider the following wholesale price vector:

$w_{L^o}^d = p_{L^o}^u$ with

$$w_{\overline{L^o}}^d = w_{\overline{L^o}}^o - \widehat{B}_{\overline{L^o}}^{-1}\left(B_{\overline{L^o}} - \widehat{B}_{\overline{L^o}}\right)\left(p_{\overline{L^o}}^o - w_{\overline{L^o}}^o\right) - \widehat{B}_{\overline{L^o}}^{-1}B_{\overline{L^o}, L^o}\left(p_{L^o}^u - w_{L^o}^o\right).$$

In addition, consider the following buy-back menu:

$$s_i^d(w_i^d, p_i) = p_i - \frac{p_i - w_i^d}{F_i(u_i^o)} \text{ for any } p_i \text{ and } i = 1, 2, ..., n$$

For any retailer i, it's profit $p_i\left(w_i^d, s_i^d(w_i^d, \bullet), p, u_i\right)$ is given by:

$$= (p_i - w_i^d)\left(a_i - \sum_{i=1}^n b_{ij}p_j\right) + (p_i - w_i^d)u_i - [p_i - s_i^d(w_i^d, p_i)]o_i$$

and is maximized when $u_i = u_i^o$ since the derivative of $(p_i - w_i^d)u_i - [p_i - s_i^d(w_i^d, p_i)]o_i$ with respect to u_i, $p_i - w_i^d - [p_i - s_i^d(w_i^d, p_i)]F_i(u_i)$ is decreasing in u_i (for feasible p_i, i.e. $p_i \geq w_i^d$) and equals to zero when $F_i(u_i) = \frac{p_i - w_i^d}{p_i - s_i^d(w_i^d, p_i)} = F_i(u_i^o)$. Hence the optimal value of u_i is u_i^o.

The partial derivative of $\Pi_i^r\left(w_i^d, s_i^d(w_i^d, \bullet), p, u_i^o\right)$ with respect to p_i for $i = 1, 2, ..., n$ is $a + u^o - v^o - Bp - \widehat{B}(p - w^d)$. Using an analogous argument as in the competing retailers system to show how we may compute the unique Nash equilibirum of retailer prices using our proposed Algorithm, and via an analogous procedure to lemma 3.25, it may shown that

165

equilibrium price vector is $\mathbf{p}^o$. Similarly, analogous to lemma 3.25, it can also be easily shown that $\mathbf{p}^o \geq \mathbf{w}^d$.

Since the retail prices and inventories are the same for both policies, it must be that:

$\Pi_e^s \left(\mathbf{w}^d, \mathbf{s}^d \left(\mathbf{w}^d, \bullet \right) \right) + \sum_{i=1}^n \Pi_i^r \left(w_i^d, s_i^d \left(w_i^d, \bullet \right), \mathbf{p}^o, u_i^o \right) = \Pi^c \left(\mathbf{p}^o, \mathbf{u}^o \right)$

$= \Pi^s \left(\mathbf{w}^o, \mathbf{s}^o \left(\mathbf{w}^o, \bullet \right) \right) + \Pi^r \left(\mathbf{w}^o, \mathbf{s}^o \left(\mathbf{w}^o, \bullet \right), \mathbf{p}^o, \mathbf{u}^o \right)$.

Since $s_i^d \left(w_i^d, p_i^o \right) o_i^o = \left[p_i^o - \frac{p_i^o - w_i^d}{F_i(u_i^o)} \right] o_i^o = p_i^o o_i^o - \left(p_i^o - w_i^d \right) v_i^o$ for $i = 1, 2, ..., n$,

$\Pi_e^s \left(\mathbf{w}^d, \mathbf{s}^d \left(\mathbf{w}^d, \bullet \right) \right) = \left(\mathbf{w}^d - \mathbf{o} \right)^t \left[\mathbf{a} - \mathbf{B} \mathbf{p}^o + \mathbf{u}^o \right] - \left[\mathbf{s}^d \left(\mathbf{w}^d, \mathbf{p}^o \right) - \mathbf{e} \right]^t \mathbf{o}^o$

$= \left(\mathbf{w}^d - \mathbf{o} \right)^t \left[\mathbf{a} - \mathbf{B} \mathbf{p}^o + \mathbf{u}^o \right] - \mathbf{p}^o \mathbf{o}^o + \left(\mathbf{p}^o - \mathbf{w}^d \right) \mathbf{v}^o$.

Since $s_i^o \left(w_i^o, p_i^o \right) o_i^o = \left[p_i^o - \frac{p_i^o - w_i^o}{F_i(u_i^o)} \right] o_i^o = p_i^o o_i^o - \left(p_i^o - w_i^o \right) v_i^o$ for $i = 1, 2, ..., n$,

$\Pi^s \left(\mathbf{w}^o, \mathbf{s}^o \left(\mathbf{w}^o, \bullet \right) \right) = \left(\mathbf{w}^o - \mathbf{o} \right)^t \left[\mathbf{a} - \mathbf{B} \mathbf{p}^o + \mathbf{u}^o \right] - \left[\mathbf{s}^o \left(\mathbf{w}^o, \mathbf{p}^o \right) - \mathbf{e} \right]^t \mathbf{o}^o$

$= \left(\mathbf{w}^o - \mathbf{o} \right)^t \left[\mathbf{a} - \mathbf{B} \mathbf{p}^o + \mathbf{u}^o \right] - \mathbf{p}^o \mathbf{o}^o + \left(\mathbf{p}^o - \mathbf{w}^o \right) \mathbf{v}^o$

$= \Pi_e^s \left(\mathbf{w}^d, \mathbf{s}^d \left(\mathbf{w}^d, \bullet \right) \right) - \left(\mathbf{w}^d - \mathbf{w}^o \right)^t \left[\mathbf{a} - \mathbf{B} \mathbf{p}^o + \mathbf{u}^o - \mathbf{v}^o \right] \leq \Pi_e^s \left(\mathbf{w}^d, \mathbf{s}^d \left(\mathbf{w}^d, \bullet \right) \right)$

since $\mathbf{w}^d \geq \mathbf{w}^o$ by definition. $\blacksquare$

Bibliography

[1] Internet ad spending is set to top '05 high. *Wall Street Journal*, 15 November:B2, 2006.

[2] Z. Abrams, A. Goel, and S. Plotkin. Set k-cover algorithms for energy efficient monitoring in wireless sensor networks. In *Proceedings of the 3^{rd} International Symposium on Information Processing in Sensor Networks*, pages 424–432, 2004.

[3] D. Acemoglu and A. Ozdaglar. Flow control, routing and performance with a for-profit service provider. Working Paper, Massachusetts Institute of Technology, 2004.

[4] M. Adler, P. B. Gibbons, and Y. Matias. Scheduling space-sharing for Internet advertising. *Journal of Scheduling*, 5:103–119, 2002.

[5] C. C. Aggarwal, J. L. Wolf, and P. S. Yu. A framework for the optimizing of WWW advertising. In *Proceedings of Trends in Distributed Systems for Electronic Commerce: International IFIP/GI Working Conference*, pages 1–10, 1998.

[6] K. Ailawadi, P. Farris, and E. Shames. Trade promotion: Essential to selling through resellers. *Sloan Management Review*, 41:83–92, 1999.

[7] A. Amiri and S. Menon. Efficient scheduling of Internet banner advertisements. *ACM Transactions on Internet Technology*, 3:334–346, 2003.

[8] A. Bar-Noy, S. Guha, J. S. Naor, and B. Schieber. Approximating the throughput of multiple machines in real-time scheduling. *SIAM Journal on Computing*, 31:331–52, 2001.

[9] F. Bernstein and A. Federgruen. Decentralized supply chains with competing retailers under demand uncertainty. *Management Science*, 51:18–29, 2005.

[10] R. C. Blattberg and S. Neslin. *Sales Promotion*. Prentice Hall, 1990.

[11] Interactive Advertising Bureau. IAB Internet advertising revenue report. 2006.

[12] G. P. Cachon. *Handbook in Operations Research and Management Science*, chapter Supply Chain Coordination with Contracts. Elsevier, 2003.

[13] G. P. Cachon and S. Netessine. *Handbook of Quantitative Supply Chain Analysis: Modeling in the eBusiness Era*, chapter Game Theory in Supply Chain Analysis. Kluwer, 2004.

[14] G. Calinescu, C. Chekuri, M. Pál, and J. Vondrák. Maximizing a submodular set function subject to a matroid constraint. In *Proceedings of the 12^{th} Conference on Integer Programming and Combinatorial Optimization*, forthcoming, 2007.

[15] C. Chekuri and S. Khanna. A polynomial time approximation scheme for the multiple knapsack problem. *SIAM Journal on Computing*, 35:713–728, 2006.

[16] C. Chekuri and A. Kumar. Maximum coverage problem with group budget constraints and applications. In *Proceedings of the 7^{th} International Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, pages 72–83, 2004.

[17] J. Chen, D. K. Friesen, and H. Zheng. Tight bound on Johnson's algorithm for maximum satisfiability. *Journal of Computer and System Sciences*, 58:622–640, 1999.

[18] J. Clifford and M. Posner. Parallel machine scheduling with high multiplicity. *Mathematical Programming*, 89:359–383, 2001.

[19] M. Conforti and G. Cornuéjols. Submodular set functions, matroids and the greedy algorithm: Tight worst-case bounds and some generalizations of the Rado-Edmonds theorem. *Discrete Applied Mathematics*, 7:251–274, 1984.

[20] S. A. Cook. The complexity of theorem proving procedures. In *Proceedings of the 3rd Annual ACM Symposium on Theory of Computing*, pages 151–158, 1971.

[21] W. J. Cook, W. H. Cunningham, W. R. Pulleyblank, and A. Schrijver. *Combinatorial Optimization*. Wiley-Interscience, 1998.

[22] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *Journal of Symbolic Computing*, 9:251–280, 1990.

[23] J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. Selfish routing in capacitated networks. *Mathematics of Operations Research*, 29:961–976, 2004.

[24] J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. On the inefficiency of equilibria in congestion games. In *Proceedings of the 11th International Conference on Integer Programming and Combinatorial Optimization*, pages 167–181, 2005.

[25] M. Dawande, J. Kalagnanam, P. Keskinocak, R. Ravi, and F. S. Salman. Approximation algorithms for the multiple knapsack problem with assignment restrictions. *Journal of Combinatorial Optimization*, 4:171–186, 2000.

[26] M. Dawande, S. Kumar, and C. Sriskandarajah. Improved performance bounds for web advertisement scheduling. In *Proceedings of the 11th Annual Workshop on Information Technologies and Systems*, pages 17–22, 2001.

[27] M. Dawande, S. Kumar, and C. Sriskandarajah. Scheduling advertisements on a web page: New and improved approximation algorithms. In *Proceedings of the Brazilian Symposium on Graphs, Algorithms and Combinatorics*, pages 59–62, 2001.

[28] M. Dawande, S. Kumar, and C. Sriskandarajah. Algorithms for scheduling advertisements on a web page: New and improved performance bounds. *Journal of Scheduling*, 6:373–393, 2003.

[29] V. Martinez de Albéniz and D. Simchi-Levi. Competition in the supply option market. Working Paper, Massachusetts Institute of Technology, 2003.

[30] B. C. Dean and M. X. Goemans. Improved approximation algorithms for minimum-space advertisement scheduling. In *Proceedings of the 30^{th} International Colloquium on Automata, Languages and Programming*, pages 1138–1152, 2003.

[31] S. Dobzinski and M. Schapira. An improved approximation algorithm for combinatorial auctions with submodular bidders. In *Proceedings of the 17^{th} Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1064–1073, 2006.

[32] J. Edmonds. Matroids and the greedy algorithm. *Mathematical Programming*, 1:127–136, 1971.

[33] H. Emmons and S. M. Gilbert. Note. The role of returns policies in pricing and inventory decisions for catalogue goods. *Management Science*, 44:276–283, 1998.

[34] A. Farahat and C. Barnhart. Improved bounds and a generalized greedy algorithm for maximizing submodular functions over matroids. Working Paper, Massachusetts Institute of Technology, 2004.

[35] U. Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45:634–652, 1998.

[36] U. Feige and J. Vondrák. Approximation algorithms for allocation problems: Improving the factor of $1 - 1/e$. In *Proceedings of the 47^{th} Annual IEEE Symposium on Foundations of Computer Science*, pages 667–676, 2006.

[37] J. Feldman, S. Muthukrishnan, M. Pál, and C. Stein. Budget optimization in search-based advertising auctions. In *Proceedings of the 8^{th} ACM Conference on Electronic Commerce*, forthcoming, 2007.

[38] M. L. Fisher, G. L. Nemhauser, and L. A. Wolsey. An analysis of approximations for maximizing submodular set functions - II. *Mathematical Programming Study*, 8:73–87, 1978.

[39] L. K. Fleischer, M. X. Goemans, V. S. Mirrokni, and M. Sviridenko. Tight approximation algorithms for maximum general assignment problems. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 611–620, 2006.

[40] A. Freund and J. S. Naor. Approximating the advertisement placement problem. In *Proceedings of the 9th International Conference on Integer Programming and Combinatorial Optimization*, pages 415–424, 2002.

[41] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.

[42] S. Fujishige. *Submodular Functions and Optimization*. Elsevier, 2005.

[43] H. Gabow and T. Kohno. A network-flow-based scheduler: Design, performance history, and experimental analysis. *Journal of Experimental Algorithmics*, 6:1–40, 2001.

[44] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1979.

[45] M. X. Goemans, V. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 142–154, 2005.

[46] D. Goldfarb and S. Liu. An $O(n^3 L)$ primal interior point algorithm for convex quadratic programming. *Mathematical Programming*, 49:325–340, 1991.

[47] R. Graham. Bounds on multiprocessing timing anomalies. *SIAM Journal on Applied Mathematics*, 17:416–429, 1969.

[48] D. Granot and S. Yin. On the effectiveness of returns policies in the price-dependent newsvendor model. *Naval Research Logistics*, 52:765–779, 2005.

[49] M. Helft. Google tests an ad idea: Pay only for results. *The New York Times*, March 21, 2007.

[50] D. S. Hochbaum, editor. *Approximation Algorithms for NP-Hard Problems*. PWS, 1997.

[51] D. S. Hochbaum and A. Pathria. Analysis of the greedy approach in problems of maximum k-coverage. *Naval Research Logistics*, 45:615–627, 1998.

[52] D. L. Hoffman and T. P. Novak. Advertising pricing models for the world wide web. In D. Hurley, B. Kahin, and H. Varian, editors, *Internet Publishing and Beyond: The Economics of Digital Information and Intellectual Property*, 2000.

[53] J. L. Holley. Note on the inversion of the Leontief matrix. *Ecnonometrica*, 19:317–320, 1951.

[54] R. A. Horn and C. R. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, 1991.

[55] Interactive advertising bureau. http://www.iab.net.

[56] O. H. Ibarra and C. E. Kim. Fast approximation algorithms for the knapsack and sum of subset problem. *Journal of the ACM*, 22:463–468, 1975.

[57] V. S. Jacob, S. Kumar, and C. Sriskandarajah. Scheduling advertising at a web site. In *Proceedings of the 10^{th} Annual Workshop on Information Technologies and Systems*, pages 37–42, 2000.

[58] K. Jansen and G. Zhang. On rectangle packing: Maximizing benefits. In *Proceedings of the 15^{th} Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 204–213, 2004.

[59] R. Johari and J. N. Tsitsiklis. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research*, 29:407–435, 2004.

[60] D. S. Johnson. Approximation algorithms for combinatorial problems. *Journal of Computer and System Sciences*, 9:256–278, 1974.

[61] H. Kellerer. A polynomial time approximation scheme for the multiple knapsack problem. In *Proceedings of the 2nd International Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, pages 51–62, 1999.

[62] H. Kellerer, U. Pferschy, and D. Pisinger. *Knapsack Problems*. Springer, 2004.

[63] S. Khot, R. J. Lipton, E. Markakis, and A. Mehta. Inapproximability results for combinatorial auctions with submodular utility functions. In *Proceedings of the 1st International Workshop on Internet and Network Economics*, pages 92–101, 2005.

[64] S. Khuller, A. Moss, and J. S. Naor. The budgeted maximum coverage problem. *Information Processing Letters*, 70:39–45, 1999.

[65] T. C. Koopmans. *Activity Analysis of Production and Allocation*. John Wiley & Sons, 1951.

[66] B. Korte and D. Hausmann. An analysis of the greedy algorithm for independence systems. In B. Alspach, P. Hell, and D. J. Miller, editors, *Algorithmic Aspects of Combinatorics*, volume 2 of *Annals of Discrete Mathematics*, pages 65–74. North-Holland, 1978.

[67] B. Korte and J. Vygen. *Combinatorial Optimization: Theory and Algorithms*. Springer, 2000.

[68] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, pages 404–413, 1999.

[69] S. Kumar, V. S. Jacob, and C. Sriskandarajah. Scheduling advertisements on a web page to maximize revenue. *European Journal of Operational Research*, 173:1067–1089, 2006.

[70] M. L. Lariviere and E. L. Porteus. Selling to the newsvendor: An analysis of price-only contracts. *Manufacturing & Service Operations Management*, 3:293–305, 2001.

[71] B. Lehmann, D. Lehmann, and N. Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55:270–296, 2006.

[72] L. Lovász. Submodular functions and convexity. In A. Bachem, M. Grötschel, and B. Korte, editors, *Mathematical Programming: The State of the Art*, pages 234–257. Springer, 1983.

[73] S. Martello and P. Toth. *Knapsack Problems: Algorithms and Computer Implementations*. John Wiley & Sons, 1990.

[74] P. Milgrom and J. Roberts. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, 58:1255–1277, 1990.

[75] V. S. Mirrokni and A. Vetta. Convergence issues in competitive games. In *Proceedings of the 7th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, pages 183–194, 2004.

[76] R. D. C. Monteiro and I. Adler. Interior path following primal-dual algorithms. *Mathematical Programming*, 44:43–66, 1989.

[77] H. Moulin. The price of anarchy of serial, average and incremental cost sharing. Working Paper, Rice University, 2005.

[78] G. L. Nemhauser and L. A. Wolsey. Best algorithms for approximating the maximum of a submodular set function. *Mathematics of Operations Research*, 3:177–188, 1978.

[79] G. L. Nemhauser and L. A. Wolsey. *Integer and Combinatorial Optimization*. Wiley-Interscience, 1988.

[80] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions - I. *Mathematical Programming*, 14:265–294, 1978.

[81] The state of online advertising, fourth quarter, 2002. Nielsen/NetRatings.

[82] C. Papadimitriou. Algorithms, games, and the Internet. In *Proceedings of the 33^{rd} Annual ACM Symposium on Theory of Computing*, pages 749–753, 2001.

[83] B. A. Pasternack. Optimal pricing and return policies for perishable commodities. *Marketing Science*, 4:166–176, 1985.

[84] G. Perakis. The price of anarchy under nonlinear and asymmetric costs. *Mathematics of Operations Research*, To appear.

[85] G. Perakis and G. Roels. The price of anarchy in supply chains: Quantifying the efficiency of price-only contracts. *Management Science*, To appear.

[86] N. C. Petruzzi and M. Dada. Pricing and the newsvendor problem: A review with extensions. *Operations Research*, 47:183–194, 1999.

[87] R. Rado. Note on independence relations. In *Proceedings of the London Mathematical Society*, volume 7, pages 300–320, 1957.

[88] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

[89] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.

[90] T. Roughgarden and É. Tardos. How bad is selfish routing ? *Journal of the ACM*, 49:236–259, 2002.

[91] T. Roughgarden and É. Tardos. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior*, 47:389–403, 2004.

[92] D. Schmeidler. Equilibrium points of nonatomic games. *Journal of Statistical Physics*, 7:295–300, 1973.

[93] A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer, 2003.

[94] D. B. Shmoys and É. Tardos. An approximation algorithm for the generalized assignment problem. *Mathematical Programming*, 62:461–474, 1993.

[95] Y. S. Song, S. Ray, and S. Li. Structural properties of buy-back contracts for price-setting newsvendors. *Manufacturing & Service Operations Management*, To appear.

[96] J. Spengler. Vertical integration and antitrust policy. *Journal of Political Economy*, 58:347–352, 1950.

[97] A. Steinberg. A strip-packing algorithm with absolute performance bound 2. *SIAM Journal on Computing*, 26:401–409, 1997.

[98] G. Strang. *Linear Algebra and its Applications*. Harcourt Jovanovich, 1988.

[99] M. Sviridenko. A note on maximizing a submodular set function subject to a knapsack constraint. *Operations Research Letters*, 32:41–43, 2004.

[100] V. Vazirani. *Approximation Algorithms*. Springer, 2001.

[101] A. Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science*, pages 416–425, 2002.

[102] X. Vives. On the efficiency of Bertrand and Cournot equilibria with product differentiation. *Journal of Economic Theory*, 36:166–175, 1985.

[103] X. Vives. Nash equilibrium with strategic complementarities. *Journal of Mathematical Economics*, 19:305–321, 1990.

[104] X. Vives. *Oligopoly Pricing*. MIT Press, 1999.

[105] R. V. Vohra and N. G. Hall. A probabilistic analysis of the maximal covering location problem. *Discrete Applied Mathematics*, 43:175–183, 1993.

[106] T. M. Whitin. Inventory control and price theory. *Management Science*, 2:61–68, 1955.

[107] L. A. Wolsey. Maximising real-valued submodular functions: Primal and dual heuristics for location problems. *Mathematics of Operations Research*, 7:410–425, 1982.