THE UNIVERSITY OF CHICAGO

DYNAMIC LEARNING IN STRATEGIC GAMES AND DUALITY IN INFINITE DIMENSIONAL OPTIMIZATION

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE UNIVERSITY OF CHICAGO
BOOTH SCHOOL OF BUSINESS
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

BY
MATTHEW STERN

CHICAGO, ILLINOIS
MARCH 2017

TABLE OF CONTENTS

LIST OF FIGURES

ABSTRACT

Optimization is the foundation for mathematical models of decision making. Many of these models involve tradeoffs between exploration, exploitation and consideration of competitors' reactions. Their analysis also relies on representing conditions for optimality. This dissertation provides new theory of how agents may behave in competitive situations including conditions for learning and strategic interactions that may inhibit learning. In the context of a single decision maker, the dissertation also provides conditions for the existence of optimal dual solutions that are easily characterized and have a meaningful economic interpretation.

The first part of the dissertation considers the impact of competition on the design and implementation of dynamic learning strategies. In monopoly pricing situations, firms should optimally vary prices to learn demand. The variation must be sufficiently high to ensure complete learning. In competitive situations, however, varying prices provides information to competitors and may reduce the value of learning. Such situations may arise in the pricing of new products such as pharmaceuticals. Chapter 1 shows that firms in competitive markets can learn efficiently by setting prices which involve adding noise to myopic estimation and best-response strategies. Chapter 2 then discusses how complete learning may not be the desired outcome when actions reveal information quickly to competitors. The chapter provides a setting where this effect can be strong enough to stop learning. Surprisingly, firms may optimally reduce any variation in prices and choose not to learn demand. The result can be that the selling firms achieve a collaborative outcome instead of a competitive equilibrium. The result has implications for policies that restrict price changes or require disclosures.

The second part of the dissertation investigates the interplay between the existence of interior points and singular dual functionals, which occur in abstract optimization. While there is growing interest in solving infinite-dimensional optimization problems, many of the intuitions and interpretations common to finite dimensions do not extend to infinite dimensions. For instance, a dual solution in finite dimensional models is represented by a vector of "dual prices" that index the primal constraints and have a natural economic interpretation. In infinite dimensions, we show that this simple dual structure, and its associated economic interpretation, may fail to hold for a broad class of problems with constraint vector spaces that are Riesz spaces (ordered vector spaces with a lattice structure) that are either σ-order complete or satisfy the projection property. In these spaces we show that the existence of interior points required by common constraint qualifications for zero duality gap (such as Slater's condition) imply the existence of *singular* dual solutions that are difficult to find and interpret. We call this phenomenon the Slater conundrum: interior points ensure zero duality gap (a desirable property), but interior points also imply the existence of singular dual solutions (an undesirable property). Riesz spaces are the most parsimonious vector-space structure sufficient to characterize the Slater conundrum. Finally, we provide sufficient conditions that "resolve" the Slater conundrum; that is, guarantee that in every solvable dual there exists an optimal dual solution that is not singular.

CHAPTER 1

INTRODUCTION

Optimization is the foundation for mathematical models of decision making. Many of these models involve tradeoffs between exploration and exploitation and consideration of competitors' reactions. Their analysis also relies on representation of conditions for optimality. This dissertation provides new theory of how agents may behave in competitive situations including conditions for learning and strategic interactions that may inhibit learning. In the context of a single decision maker, the dissertation also provides conditions for the existence of optimal dual solutions that are easily characterized and have a meaningful economic interpretation.

The beginning of this dissertation focuses on developing effective pricing strategies in new or uncertain markets. Such strategies involve balancing the tradeoff between learning about consumer demand and earning profits in each period. For markets with a small number of firms, this tradeoff becomes increasingly difficult to manage as firms must not only account for the unknown demand from consumers, but also for the response from their competitors. As the market matures and firms learn from the experience of selling their products, each firm becomes more proficient at setting prices; however, the resulting increased level of competition may cause profits to decline for the industry as a whole. This effect brings into question whether it is in the best interest of firms to actively learn about their business environment, or rely only on their current information to compete for profits. In other words, do competing firms always want to learn about demand?

Our primary research interest concerns the impact of competition on the design and implementation of dynamic pricing strategies. Dynamic pricing strategies are a useful tool for firms that face a changing marketplace for their product. As market conditions evolve over time, firms can employ such strategies by adjusting their prices periodically to match their environment. For example, consider the market for a single, new product where the demand for that product is initially unknown. From the firm's perspective, the marketplace is changing as data from the sales of the product is collected, providing more information about the underlying demand. In each period, the firm must weigh the expected revenue that their pricing decision will yield in the current period against the information value that charging varied or experimental prices may provide for future pricing decisions. Balancing this learning and earning tradeoff for a monopolist firm has been the focus of recent research in revenue management and has produced effective pricing policies under various forms of uncertainty. Our goal is to assess the performance of these dynamic strategies when multiple firms are selling products in an uncertain marketplace and examine the effects of competition on the learning and earning tradeoff.

Designing strategies that account for both demand learning and the impact of competition leads to the following questions. If firms adopt dynamic pricing strategies taken from the monopolist literature for demand learning, will they eventually learn the underlying demand conditions for their products? At same time, will revealing the market conditions increase the degree of competition and "learn" away profits for all of the firms involved? To address these questions, we introduce two models of competition under demand uncertainty. The first tests the performance of various monopolist strategies and heuristics in these competitive environments, while the second explores potential equilibrium strategies for these markets.

In Chapter 1, each competing firm is allowed to choose a strategy from a family of pricing policies taken from the monopolist literature on dynamic pricing. When these strategies are applied simultaneously, we show that the prices of the individual firms can approach an overall Nash equilibrium of the game with complete information, provided the firms sufficiently vary their prices. On the other hand, if firms do not vary prices sufficiently, the firms can incompletely learn the demand and prices converge to a different set of prices other than the complete-information Nash equilibrium. These prices can be beneficial or harmful to individual firms relative to the complete-information Nash equilibrium, and, in some cases, can actually be better than the complete-information solution for all firms (an aspect which we explore more fully in the second chapter of the dissertation).

Chapter 2 addresses the issue of incomplete learning under competition by directly investigating equilibrium pricing policies in dynamic games of incomplete information. The model shows that incomplete learning of demand information is not merely a byproduct of the dynamic pricing strategies investigated in the first model, but is rather a rational outcome for firms competing in a Markov perfect equilibrium (MPE). In particular, we develop several simple demand environments, where the equilibrium strategies for the competing firms actively avoid learning the true value of demand and attain a collusive outcome even in finite time horizons. Hence, firms prefer to remain willfully ignorant of the marketplace for their products.

Chapter 3 investigates the interplay between the existence of interior points and singular dual functionals, which occur in abstract optimization. We show that the existence of interior points required by common constraint qualifications for zero duality gap (such as Slater's condition) imply that singular dual solutions exist. We call this phenomenon the Slater conundrum: interior points ensure zero duality gap (a desirable property), but interior points also imply the existence of singular dual solutions (an undesirable property). Riesz spaces are the most parsimonious vector-space structure sufficient to characterize the Slater conundrum. Finally, we provide sufficient conditions that "resolve" the Slater conundrum; that is, guarantee that in every solvable dual there exists an optimal dual solution that is not singular.

CHAPTER 2

SEMI-MYOPIC COMPETITION USING OLS ESTIMATION

Pricing when demand is uncertain requires balancing the trade-off between learning about demand and earning revenues in the short run. In order to accurately estimate the parameters of a demand model, firms must vary the prices of their products across the selling season; however, this practice typically requires firms to deviate from the myopic price that earns the most expected revenue given their current information and forecasts. In this chapter, we evaluate the performance of various myopic and semi-myopic heuristics from the revenue management literature on dynamic pricing.

Each firm is allowed to choose a strategy from a family of policies originally designed for monopolist pricing. These strategies are extended to explicitly incorporate the competitive environment by forecasting the prices from rival firms, using on past observations. We show that the prices of the individual firms can approach an overall Nash equilibrium of the game with complete information, provided the firms sufficiently vary their prices. On the other hand, if firms do not vary prices sufficiently, the firms can incompletely learn the demand and prices converge to a different set of prices rather than the complete-information Nash equilibrium. These prices can be beneficial or harmful to individual firms relative to the complete-information Nash equilibrium, and, in some cases, can actually be better than the complete-information solution for all firms.

2.1 Literature Review

The development of dynamic pricing strategies has been studied in fields varying from operations management and economics to statistics and machine learning. In this chapter, we explore the relationship between the equilibrium strategies of competing firms and the learning and earning strategies developed in the revenue management literature. This review outlines the relevant economics literature on collusive pricing and equilibrium price experimentation as well as the recent research on dynamic pricing and learning in stationary demand environments. For a comprehensive study of the dynamic pricing literature on demand learning, we refer the reader to the recent surveys of den Boer [2015] and Bitran and Caldentey [2003]. Similarly, the books of Fudenberg and Levine [1998], Kirman and Salmon [1995] and Young [2004] provide a rigorous foundation for the economics theory of learning in games.

Early works of Rothschild [1974], McLennan [1984], and Lai and Robbins [1982] showed that following a myopic strategy can lead to incomplete learning in the long run. In the case of a monopolist firm, such strategies will inevitably underperform more forward looking policies that explore the potential pricing alternatives. By introducing a performance metric known as *regret*, researchers in revenue management have developed semi-myopic policies to avoid incomplete learning by strategically experimenting with prices. However, learning the demand curve in the long run may not be a desirable strategy for firms in a competitive marketplace. For instance, firms would prefer to have mistaken beliefs if the resulting market prices outperform the full-information competitive equilibrium. Adhering to a myopic strategy would therefore be rational in competitive environments provided that firms can recognize situations where their mistaken beliefs generate collusive prices. Our work extends the notion of regret to competitive environments, develops pricing strategies that avoid incomplete learning, and analyzes the connection between demand uncertainty and collusive pricing.

Since the seminal work of Stigler [1964], market conditions that support collusive pricing policies have been well studied by the field of economics. In particular, demand uncertainty is classically seen as an obstacle to collusive behavior. When firms cannot accurately monitor the actions of their competitors, they cannot distinguish between a realized period of low demand and a rival's deviation from the collusive agreement. As firms can neither detect nor punish these individual deviations, they engage in “secret price cutting” strategies, shaving the price to the competitive equilibrium. Demand uncertainty, however, does not preclude collusion as Green and Porter [1984], Williams [2001] and Janjgava and Slobodyan [2011] show that collusive agreements can be sustained periodically over a business cycle.

Green and Porter [1984] determined that firms facing uncertain demand can alternate between periods of collusion and periods of competition. In their setting, firms compete on the basis of quantity and though each firm's action is private, the resulting market price is publicly revealed in each period. Using the market price as a noisy signal, a price war is triggered every time the market price falls below a specified level and firms switch from producing collusive quantities to producing competitive quantities. Williams [2001] and Janjgava and Slobodyan [2011] consider models of price competition and quantity competition, respectively, where two firms sequentially estimate their unknown demand curves without accounting for their opponent's actions. In contrast to the price wars characterized by Green and Porter, these misspecified models conclude that firms will learn to play the single-period competitive equilibrium but periodically “escape” to collusive levels as a result of large deviations in the realized demands or market prices. Our research builds on their approaches by allowing firms to estimate not only their unknown demand curves, but also the strategic response from their competitors in future periods. We show that the value of information for the competing firms may be negative, and firms can strategically slow the learning process in order to sustain collusive prices. Our model therefore demonstrates that uncertainty about the demand curve can enable firms to collude, rather than discourage these agreements.

A large literature in economics focuses on equilibrium price dispersion when firms learn about the degree of differentiation between their products. Aghion et al. [1993] and Harrington [1995] each consider duopoly models where firms face symmetric yet uncertain business environments. Aghion et al. [1993] consider a finite-horizon Hotelling model where each firm incurs a fixed but unknown transportation cost and Harrington [1995] examines a two-period model where each firm receives an i.i.d cost shock before setting prices. In each model, the authors characterize the subgame perfect equilibrium strategies and determine that firms will experiment with prices in equilibrium. In particular, Harrington [1995] concludes that the value of information is negative when products are highly differentiated and positive when they are close substitutes. Keller and Rady [2003] demonstrate this effect using a continuous time model where the unknown degree of product differentiation periodically fluctuates between two known values. As firms track the changes in this parameter, they alternate between periods of price experimentation and price matching depending on the current value of information. These models therefore support the claim that price experimentation is rational for models of competition with demand learning and emphasize the importance of the value of information.

The economics literature on collusion and price experimentation demonstrates the need for competing firms to take an active role in learning the demand curves for their products. However, the application of these models to business practice is limited by the fact that the equilibrium behavior for firms with incomplete information requires each firm to estimate the information available to its competitors. Blume and Easley [1995] point out that this requires "superhuman rationality" on the part of firms and restricts the complexity of the models considered in the literature. In contrast, research in revenue management focuses instead on developing implementable policies to provide tactical guidance to firms regularly faced with these business environments.

While the revenue management literature on dynamic pricing with demand learning has grown rapidly in the last decade, many of the models studied in this area focus on monopolist firms. Such models assume an underlying parametric model for demand, which a firm sequentially estimates by observing the purchase decisions of customers. The literature can be divided into two general streams of research, characterized by the statistical methods employed by the firm. Aviv and Pazgal [2005], Araman and Caldentey [2009], Farias and

Van Roy [2010], and Harrison et al. [2012] assume that a monopolist firm faces a Bayesian uncertainty about its demand environment and develop well-performing polices that balance the tradeoffs between learning and earning in these settings. Alternatively, research adopting the frequentist perspective assumes that the firm applies classical estimation techniques such as ordinary least squares or maximum-likelihood estimation to provide point estimates for the parameters in each period. The firm then faces the decision of charging the myopic or certainty-equivalent price that maximizes expected revenue according these estimates or deviating from this price to improve the quality of the estimates themselves.

In principle, the frequentist model of the monopolist pricing problem can be formulated as a dynamic program and solved using standard methods; however, even simple formulations of these learning models face the curses of dimensionality. As such, several authors have explored approximate dynamic programming methods for determining the optimal pricing strategy. Lobo and Boyd [2003] linearize the inverse of the covariance matrix for the unknown parameters about the myopic policy with "dithering." Their resulting approximation reduces the dynamic programming formulation to a solvable quadratic program with semidefinite constraints. Carvalho and Puterman [2005] consider one-step look ahead policies that choose prices to optimize the combined expected revenue of the next two periods and Bertsimas and Perakis [2006] consider a reduced state-space approximation to the dynamic program. One of the few revenue management articles to consider competition, Bertsimas and Perakis [2006] show that their approach extends to a duopoly model, where one firm estimates the parameters used by its competitor through inverse optimization.

An alternative approach to dynamic pricing with frequentist demand learning focuses on the design of well-performing, semi-myopic strategies that are simple to implement in practice. The performance of these dynamic pricing strategies is measured in terms of *regret*. For a monopolist firm, regret is the expected revenue loss incurred in T periods relative to the benchmark performance that a clairvoyant firm would enjoy given complete information of the demand parameters. Regret minimizing policies have been a vital benchmark for analyzing the tradeoffs between learning and earning for a variety of demand models and applications. Besbes and Zeevi [2009], Besbes and Zeevi [2011], Broder and Rusmevichientong [2012], and den Boer and Zwart [2013] each consider monopolist firms that sell a single product to consumers. They determine lower bounds on the asymptotic growth rate of regret in their demand environments and develop policies that achieve similar growth rates. Recently, den Boer [2014] and Keskin and Zeevi [2014] analyzed minimal regret policies for monopolists offering multiple products.

One of the few exceptions to the assumptions of monopolistic pricing in the revenue management literature is Cooper et al. [2015], which considers the impact of firms' assuming monopolistic conditions when they are actually in competition. They show that this can result in a Nash equilibrium, a cooperative solution that maximizes total producer surplus, or intermediate outcomes. The results in this chapter show that these outcomes are also possible when the agents are aware of their competition but vary the extent of their price experimentation. The results show, however, that a Nash equilibrium is always attained if each agent follows an optimal learning strategy as in a monopolistic setting.

A different literature on competitive situations with learning is from the machine learning subfield of computer science. This field focuses on the analysis of heuristic policies that decompose by agent and converge to a single-stage Nash equilibrium. For a comprehensive discussion on these policies in the context of game theory, we refer to the books of Young [2004] and Fudenberg and Levine [1998]. Other references include Schipper [2015], Foster and Young [2003], Foster and Young [2006], Hart et al. [2006], Jordan [1991], and Kalai and Lehrer [1993].

2.2 Model

Consider N firms, each about to introduce a single, new product to the marketplace. Firms are aware of their competitors' products and that the prices chosen by their competitors will influence the demand for their own product. The underlying market conditions, however, are initially unknown. The demand for each distinct product is related to the others, but it is uncertain whether these products will behave as substitutes or complements, nor is it clear what magnitude a change each product's own price will have on its demand. Firms learn about the market conditions by pricing their product and privately observing their resulting sales.

We model their interactions as a T-period, dynamic game of incomplete information. Let $I := \{1, \ldots, N\}$ be the set of firms. The market conditions are characterized by a state of the world, which remains constant throughout the selling season, and is drawn randomly before selling begins from a commonly known set of demand parameters. In each period $t \in \{1, \ldots, T\}$, firms price their products publicly and simultaneously with firm $i \in I$ choosing its price p_{it} from a closed interval of the positive real line $P \subset \Re^{+}$. After prices are chosen, each firm then receives a private signal $d_{it} \in \Re$ drawn randomly from consumers. For our purposes, each firm's private signal represents that firm's observed demand for their product in period t. Firms have no capacity limitations for their products and we ignore the effects of marginal and fixed costs to focus instead on the demand learning and earning tradeoff.

Drawing from the setup in Keskin and Zeevi [2014], each firm correctly assumes that their demand is a linear function of the prices charged in each period,

$$d_{it} = \theta_i^\top \begin{bmatrix} 1 \\ p_{1t} \\ \vdots \\ p_{Nt} \end{bmatrix} + \epsilon_{it}, \quad i = 1, \ldots, N.$$

The unknown market conditions $\theta = (\theta_1, \ldots, \theta_N)^\top$ are initially drawn from a compact set $\Theta \subset \Re^{N(N+1)}$ and the demand shocks, ϵ_{it} are drawn from a static, mean-zero distribution independently for each of the N products and serially independent across time. For convenience, let $x_t = (1,\ p_t^\top)^\top$, where $p_t = (p_{1t}, \ldots, p_{Nt})^\top$ is the vector of prices chosen by each firm in period t. It will also be useful to distinguish between the various demand parameters within each firm's demand model. Thus, with a slight overuse of notation, let $\theta_{i0} = \alpha_i$, $\theta_{ii} = \beta_i$ and $\theta_{ij} = \gamma_{ij}$ for $i \neq j$. The demand curve for firm i can now be expressed as

$$d_{it} = \alpha_i + \beta_i p_{it} + \sum_{j \neq i} \gamma_{ij} p_{jt} + \epsilon_{it}, \quad i = 1, \ldots, N.$$

Though equivalent, both the vector and component form for the demand curves will be useful for our analysis as will be evident in context throughout the chapter.

At the start of each period t, firms use their available data to estimate the state of the world, to forecast the prices that their competitors will choose, and to determine their own pricing action. We describe each agent's strategy as comprising these elements of estimation, forecast, and action, all based on public information revealed in the past actions of all agents and each individual agent's private information on past payoffs.

In considering demand estimates, we note that observed sales represent samples of the uncensored demand distribution since we have ignored capacity considerations. As a simple estimation strategy, we consider that each agent uses ordinary least squares regression. At the start of period t, firm i's least squares estimator of θ_i is

$$\hat{\theta}_{it} = \left(\sum_{s=1}^{t-1} x_s x_s^\top\right)^{-1} \left(\sum_{s=1}^{t-1} d_{is} x_s\right).$$

Since each firm's realized demand history $\{d_{is}\}_{i=\pm1,s=1,\ldots,t}$ is private information, only firm i has knowledge of the least-squares estimate $\hat{\theta}_{it}$. However, the least-squares estimates for the N firms are interconnected as they share the same empirical Fisher information matrix given by

$$\mathcal{J}_t \;=\; \sum_{s=1}^{t} x_s x_s^\top.$$

This matrix plays a crucial role in our analysis, as its size (as measured by its smallest eigenvalue $\lambda_{\min}(\mathcal{J}_t) \geq 0$) indicates how close each firm's estimates are to the true value of the parameters. Combining the definition of the least-squares estimates with the true demand model yields the following expression for the estimation error:

$$\hat{\theta}_{it} - \theta_i = \mathcal{J}_t^{-1} \sum_{s=1}^{t} \epsilon_{is} x_t.$$

As the estimation error is inversely proportional to the empirical Fisher information, increasing the size of this matrix will enable each firm to learn their underlying demand; however, joint control of the Fisher information is non-trivial, since it is generated through the combined pricing actions of all of the firms in the market. The distribution of the demand noise ϵ_{it} plays a important role in the estimation error as well. To ensure that the demand noise does not dominate the estimation error, we assume that it follows a light-tailed distribution, as in Keskin and Zeevi [2014]; that is, there exists a positive constant z_0 such that $\mathbb{E}[e^{z\epsilon_{it}}] < \infty$ for all $|z| < z_0$. The least-squares estimates can be improved further by using the knowledge that the true parameters θ belong to the compact set Θ. Let

$\vartheta_{it} := \text{argmin}_{\vartheta \in \Theta_i}\{||\vartheta - \hat{\theta}_{it}||\}$ be the $L2-$projection of the least-squares estimates onto firm i's subset of the demand parameter set $\Theta_i \subset \Theta$.

Along with their demand estimates, each firm develops a forecast for their competitors' future prices p^e_t at start of period t. We allow a range of assumptions that have appeared in the literature:

1. *Cournot adjustment*: $p^e_t = p_{t-1}$;
2. *Time average over H-horizon:* $p^e_t = \frac{1}{H}\sum_{\tau=t-H}^{t-1} p_\tau$ for $H < T$;
3. *Exponential smoothing:* $p^e_t = \sum_{\tau=1}^{t-1} \lambda^{t-\tau} p_\tau$ for $\lambda < 1$.

Each firm selects one of these forecasting strategies before selling begins; indeed, rival firms need not be aware of their competitors' choices. Given these estimates and forecasts, we define an estimated expected demand for each agent i's demand given their action p_{it} as $\hat{d}_{it}(p^e_t, p_{it}, \vartheta_{it}))$.

A general pricing strategy for firm i is a sequence of functions $\sigma_i := (\sigma_{i1}, \sigma_{i2}, \dots \sigma_{iT})$ where $\sigma_{it} : \Re^{(N+1)(t-1)} \to \mathcal{P}(P)$ is a measurable mapping from firm i's observable history $\mathcal{H}_{it} = (d_{i1}, p_1, \dots, d_{i(t-1)}, p_{(t-1)})$, to the space of probability measures on the closed interval of prices $P \subset \Re^+$. However, by introducing the demand estimates and price forecasts, we consider a restricted set of *admissible* pricing strategies $\sigma_{it} : \Re^{2N+1} \to \mathcal{P}(P)$ such that firm i's price in period t is distributed as

$$p_{it} \sim \sigma_{it}(\vartheta_{it}, p^e_t).$$

Our analysis revolves around strategies that focus on best-response (as, for example, also assumed in Simon [2007]). Consider the firm's expected single-stage payoffs given their estimates and forecasts. In this case, firms can maximize their single-stage profits by choosing prices that maximize expected revenues (since we assume zero marginal cost). Each firm's profit in a given period t is their best-response to the forecast prices of competitors defined as the product of their chosen price and the consumer demand $u_i(p_t, d_{it}) = p_{it}d_{it}$.

The single-period best-response function for firm i given an estimate ϑ_{it} and price forecast p_t^e is

$$\varphi(\vartheta_{it}, p_t^e) = \arg\max_{p_{it}} u_i(p_{it}, \hat{d}_{it}(p_t^e, p_{it}, \vartheta_{it})) = \frac{\hat{\alpha}_i}{-2\hat{\beta}_i} + \sum_{j \neq i} \frac{\hat{\gamma}_{ij}}{-2\hat{\beta}_i} p_{jt}^e.$$

If the true demand parameters θ were known, then solving for the best-response prices for each firm simultaneously yields the unique single-period Nash equilibrium prices

$$p_i^{NE} = \varphi(\theta_i, p^{NE}).$$

Since the best-response in a monopolistic setting can yield incomplete learning, we add a perturbation to the best-response in the form of a noise term. We call this set of policies *best response with random dithering* (following Lobo and Boyd [2003]), given as follows:

$$\sigma_{it}(\vartheta_i, p_t^e) = \varphi(\vartheta_{it}, p_t^e) + \nu_{it},$$

such that the ν_{it} are mean-zero random variables, that are conditionally independent over i and t given the history at time t. Therefore, at the start of period t and given history $\mathcal{H}_{it}$, firm i's price is a random variable with distribution $\sigma_{it}(\vartheta_i, p_t^e)$ and mean $\varphi(\vartheta_{it}, p_t^e)$. A simple way to generate admissible random dithering policies is to choose distributions for the random noise terms before the selling season begins and then truncate these distributions in each period so that prices remain in P.

Note that in order for best response with random dithering polices to be *admissible* strategies, the prices that each firm charges in each period must lie in a closed interval P. An immediate consequence of this property is that the ν_{it} noise terms must be bounded in each period and that the best response function $\varphi(\vartheta_i, p^e)) \in P$ for all $\vartheta_i \in \Theta_i$ and all $p^e \in P^N$. To that end, we assume that each firm's best-response price lies within an open interval $\text{int}(P)$ whenever its competitors' prices also lie within that interval. This assumption allows for the added noise terms ν_{it} to have positive variance in each period and has the following implications for the parameter space Θ.

Proposition 2.2.1. *If $\varphi(\vartheta_i, p_t^e)) \in \text{int}(P)$ for all $\vartheta_i \in \Theta_i$ and all $p_t^e \in P^N$ then,*

i. $\beta_i < 0$ *for all* $i \in 1 \ldots N$ *and all* $\theta \in \Theta$;

ii. $|\sum_{j \neq i} \gamma_{ij}| \leq -2\beta_i$ *for all* $i \in 1 \ldots N$ *and all* $\theta \in \Theta$.

Proof. First note that if $\beta_i \geq 0$, then demand for product i increases with the price of product i and by definition the best response would be unbounded $\varphi(\vartheta_i, p_t^e) = \infty$. Let $P = [l, u]$ where $u > l > 0$. Then the second property comes from enforcing that $\varphi(\vartheta_i, ue) \leq u$ and $\varphi(\vartheta_i, le) \geq l$, where e is the N-vector of ones. Combining these inequalities yields $\varphi(\vartheta_i, ue) - \varphi(\vartheta_i, le) \leq u - l$ or $(\sum_{j \neq i} \hat{\gamma}_{ij} / -2\hat{\beta}_i)(u - l) \leq u - l$. Since both $-2\hat{\beta}_i$ and $u - l$ are positive, this implies that $\sum_{j \neq i} \hat{\gamma}_{ij} \leq -2\hat{\beta}_i$. Similarly, the inequality $-(\sum_{j \neq i} \hat{\gamma}_{ij}) \leq -2\hat{\beta}_i$ follows through the combination of $\varphi(\vartheta_i, ue) \geq l$ and $\varphi(\vartheta_i, le) \leq u$. □

Additionally, it is necessary to strengthen the second property of Proposition 2.2.1 and assume that the parameter space has the property that $|\sum_{j \neq i} \gamma_{ij}| \leq -\beta_i$ for all $i \in 1 \ldots N$ and all $\theta \in \Theta$. This assumption is required due to a byproduct of our proof technique which bounds the forecasting errors inherent to our allowed options for competitor forecasts.

We restrict the analysis below to this limited set of admissible policies since finding equilibria in fully general settings is not realistic. It requires, for example, extraordinary rationality (see Blume and Easley [1995]) for each agent to fully consider all other agents' information states, updating capabilities, and strategy choices. In the following chapter, however, we analyze a restricted action and parameter space setting where we can more fully describe equilibria in the repeated game.

2.3 Measuring Regret

In order to analyze the pricing strategies outlined in the previous section, we introduce a performance metric called *regret*, which compares the revenues generated by each firm to the revenues generated by the competitive Nash equilibrium. The T-period regret for firm i is given by

$$\Delta_i^\sigma(\theta, T) = \sum_{t=1}^{T} \mathbb{E}_\theta^\sigma[p_i^{NE}(\alpha_i + \beta_i p_i^{NE} + \sum_{j \neq i} \gamma_{ij} p_j^{NE}) - p_{it}(\alpha_i + \beta_i p_{it} + \sum_{j \neq i} \gamma_{ij} p_{jt})],$$

where the expectation $\mathbb{E}_\theta^\sigma$ is taken with respect to the probability of the price sequences $\{p_t\}_{t=1,\ldots,T}$ induced by the firms' chosen strategies σ and the true value of demand parameters θ. We note that this is distinct from the definition of regret often used in the machine learning literature (see Zinkevich et al. [2007]). In that context, regret is measured as the difference relative to the best possible static competitor *actions*; that is, rather than comparing realized profits to the complete-information Nash equilibrium, firms would instead compare their outcomes to the profits they would have accrued if their competitors chose to charge a constant price in each period. We consider our admissible policy responses as more fully capturing rational behavior. As all firms are pricing to learn demand and earn revenues, charging a constant price would prohibit learning not only for those firms that choose to adopt that strategy, but also for the marketplace as a whole.

Choosing to benchmark against the competitive Nash outcome is also useful because it has a direct analogue to the monopolist learning and earning literature. In particular, it separates into two quantities of interest that we describe as the regret due to learning and the regret due to influence.

$$\Delta_i^\sigma(\theta, T) \quad = \quad -\beta_i \sum_{t=1}^{T} \mathbb{E}_\theta^\sigma\left[(p_i^{NE} - p_{it})^2\right] - 2\beta_i \sum_{t=1}^{T} \mathbb{E}_\theta^\sigma\left[p_{it}\left(\varphi(\theta_i, p^{NE}) - \varphi(\theta_i, p_t)\right)\right].$$

The first term in the sum measures the expected squared distance between firm i's Nash equilibrium price and the prices chosen in each period according to its strategy σ_i. We identify this term as measuring firm i's regret due to learning, as it tracks the firm's knowledge of the underlying demand parameters and the systems progress towards the complete-information Nash equilibrium. The expected squared distance is precisely the regret metric that a monopolist would face in this market when balancing the tradeoffs between learning and earning. The worst-case regret of a monopolist is

$$\Delta^\sigma(T) \leq \sup\{-|K| \sum_{t=1}^{T} \mathbb{E}_\theta^\sigma \left[||p^{CE} - p_t||^2 \right] : \theta \in \Theta\},$$

where p^{CE} is the vector of cooperative Nash equilibrium prices and K is a known constant. The remaining term, called the *regret due to influence*, represents the regret that firm i incurs as a result of the mistaken beliefs about its competitor. Notice that when the products in the market are substitutes (complements), this term is negative and when competing firms charges above (below) its Nash equilibrium price. Our definition of regret allows for such negative values as in learning scenarios, a firm i can exploit the ignorance of its competitors to earn additional revenues. Hence, firms may consider allowing an increase in the regret due to learning in an effort to influence their competitor to charge a more favorable price.

2.4 Conditions for Efficient Learning

This section analyzes the use of the admissible strategies in achieving an efficient learning outcome, as measured by the worst-case regret due to learning. In particular, we show first that the information grows at the rate of the variance of the pricing strategies σ, as controlled by the noise terms in random dithering policies. These results depend on the information metric $\mathcal{J}_t$ and its minimum eigenvalue, denoted $\lambda_{\min}(\mathcal{J}_t)$. To bound this minimum eigenvalue, we use the following matrix version of the Freedman bound. Let $\mathbb{E}_s$

and Var_s denote the conditional expectation and variance of an adapted sequence of random matrices given a history of realizations up to time s.

Theorem 2.4.1 (Matrix Freedman, Tropp et al. [2011])**.** *Consider a finite adapted sequence $\{Y_s\}$ of random, self-adjoint matrices with dimension d. Assume that*

$$\mathbb{E}_{s-1}Y_s = 0 \text{ and } \lambda_{\max}(Y_s) \leq R \text{ almost surely.}$$

Define the finite series

$$Z := \sum_s Y_s \text{ and } W := \sum_s \mathbb{E}_{s-1}(Y_s^2).$$

Then for $\delta \geq 0$ and $v > 0$,

$$\mathbb{P}\left\{\lambda_{\max}(Z) \geq \delta \text{ and } \lambda_{\max}(W) \leq v^2\right\} \leq d \cdot \exp\left(\frac{-\delta^2/2}{v^2 + R\delta/3}\right).$$

The Freedman bound can be applied directly to our model as follows:

Corollary 2.4.2. Let $\mathcal{J}_t = \sum_{s=1}^{t} x_s x_s^\top$ be the Fisher information matrix at time $t \leq T$ generated by admissible strategies σ. Assume that $\exists v > 0$ such that

$$||\sum_{s=1}^{t} \mathrm{Var}_{s-1}[x_s x_s^\top]|| < v, \qquad \text{almost surely}$$

then

$$\mathbb{P}\left\{\lambda_{\min}(\mathcal{J}_t) \leq -\delta + \sum_{s=1}^{t} \lambda_{\min}\left(\mathbb{E}_{s-1}[x_s x_s^\top]\right)\right\} \leq N \cdot \exp\left(\frac{-\delta^2/2}{v^2 + R\delta/3}\right).$$

Proof. This follows through a direct application of the Freedman bound and Weyl's inequalities. Since the sequence of public price vectors x_s is a finite sequence of adapted random vectors, the random matrices $Y_s := \mathbb{E}_{s-1}[x_s x_s^\top] - x_s x_s^\top$ form an adapted sequence of random, self-adjoint matrices each with conditional mean zero.

By the definition of our admissible strategies, the prices charged in each period lie in an interval $P = [l, u]$ with $0 < l < u$, so $||x_t||^2 \leq Nu^2 := R$. The fact that prices are uniformly bounded implies that the maximum eigenvalues of the Y_s are also uniformly bounded by the following argument. Consider the spectral norm of the Y_s,

$$||Y_s|| = ||\mathbb{E}_{s-1}[x_s x_s^\top] - x_s x_s^\top|| = \max\{||\mathbb{E}_{s-1}[x_s x_s^\top]||, ||x_s x_s^\top||\} \leq R.$$

The above expression for spectral norm follows from the fact that both $\mathbb{E}_{s-1}[x_s x_s^\top]$ and $x_s x_s^\top$ are positive semi-definite matrices and the bound follows from the following application of Jensen's inequality:

$$||\mathbb{E}_{s-1}[x_s x_s^\top]|| \leq \mathbb{E}_{s-1}||x_s x_s^\top|| \leq \mathbb{E}_{s-1}||x_s||^2 \leq R.$$

Since the spectral norm of a Hermitian matrix is $||X|| = \max\{\lambda_{\max}(X), -\lambda_{\min}(X)\}$, we have that

$$\lambda_{\max}(Y_s) \leq ||Y_s|| \leq R.$$

Applying the Matrix Freedman bound yields:

$$\begin{aligned}
n \exp\left(\frac{-\delta^2/2}{\sigma^2 + R\delta/3}\right) &\geq \mathbb{P}\left\{\lambda_{\max}\left(-\mathcal{J}_t + \sum_{s=1}^{t} \mathbb{E}_{s-1}[x_s x_s^\top]\right) \geq \delta\right\} \\
&\geq \mathbb{P}\left\{\lambda_{\max}(-\mathcal{J}_t) \geq \delta - \sum_{s=1}^{t} \lambda_{\min}(\mathbb{E}_{s-1}[x_s x_s^\top])\right\} \\
&= \mathbb{P}\left\{\lambda_{\min}(\mathcal{J}_t) \leq -\delta + \sum_{s=1}^{t} \lambda_{\min}(\mathbb{E}_{s-1}[x_s x_s^\top])\right\}.
\end{aligned}$$

Notice that the matrix $W = \sum_{s=1}^{t} \mathrm{Var}_{s-1}[x_s x_s^\top]$ in the Freedman bound is omitted. This is due to the fact that the spectral norm of a matrix is greater than or equal its maximum eigenvalue and the condition $\lambda_{\max}(W) \leq v^2$ holds almost surely by assumption. $\square$

We can now use the corollary above to bound the minimum eigenvalues of the Fisher information matrix in our setting.

Lemma 2.4.3. If firms choose admissible strategies and there exist constants $c_i, C_i > 0$ such that $\frac{c_i}{\sqrt{t}} \leq \text{Var}_{t-1}[\sigma_{it}] \leq \frac{C_i}{\sqrt{t}}$ almost surely for $i = 1, \ldots, N$ and $t < T$ then there exist constants $\kappa_0, \kappa_1 > 0$ such that

$$\mathbb{P}(\lambda_{\min}(\mathcal{J}_t) < \kappa_0 \sqrt{t}) \leq \frac{\kappa_1}{\sqrt{t}}.$$

Proof. The proof will proceed as follows. First, we will provide a lower bound for the minimum eigenvalue for the sum of the conditional expectations of $x_s x_s^\top$ for $s = 1, \ldots, t$. Second, we will provide an upper bound for the sum of the conditional variances of $x_s x_s^\top$ for $s = 1, \ldots, t$. These bounds then allow us to apply the Matrix Freedman corollary and complete the proof. Express the price vectors in each period as the sum

$$x_s = z_s + \nu_s.$$

where $z_s := \mathbb{E}_{s-1}[x_s]$ and $\nu_s \in \Re^{N+1}$ is a random vector with zero mean and variance equal to $\text{Var}_{s-1}[x_s]$.

Claim 1. $\lambda_{\min}\left(\sum_{s=1}^{t} \mathbb{E}_{s-1}[x_s x_s^\top]\right) \geq \delta_0 \sqrt{t}$ *for some* $\delta_0 > 0$.

Proof of Claim 1

Separate the conditional expectations into z_s and ν_s components,

$$\begin{aligned}
\sum_{s=1}^{t} \mathbb{E}_{s-1}[x_s x_s^\top] &= \sum_{s=1}^{t} z_s z_s^\top + \sum_{s=1}^{t} \mathbb{E}_{s-1}[\nu_s \nu_s^\top] \\
&= \sum_{s=1}^{t} (z_s - \bar{z})(z_s - \bar{z})^\top + \sum_{s=1}^{t} \bar{z}\bar{z}^\top + \sum_{s=1}^{t} \mathbb{E}_{s-1}[\nu_s \nu_s^\top] \\
&\succeq \sum_{s=1}^{t} \bar{z}\bar{z}^\top + \sum_{s=1}^{t} \mathbb{E}_{s-1}[\nu_s \nu_s^\top],
\end{aligned}$$

where $\bar{z} := \frac{1}{t}\sum_{s=1}^{t} z_s$. Let $y = (y_1, \ldots, y_{N+1}) \in \Re^{N+1}$ be an arbitrary unit vector and y_* and $\bar{z}_*$ be vectors consisting of the last N components of $y, \bar{z}$, respectively. Let $c = \min_{i=1,\ldots N}\{c_i\}$, then

$$\begin{aligned}
y^\top \left(\sum_{s=1}^{t} \bar{z}\bar{z}^\top + \sum_{s=1}^{t} \mathbb{E}_{s-1}[\nu_s \nu_s^\top] \right) y &= \sum_{s=1}^{t} (y_1 + y_*^\top \bar{z}_*)^2 + \sum_{i=1}^{N} y_{i+1}^2 \mathrm{Var}_{s-1}[\sigma_{is}] \\
&\geq \sum_{s=1}^{t} (y_1 - ||y_*|| \cdot ||\bar{z}_*||)^2 + ||y_*||^2 \frac{c}{\sqrt{t}} \\
&\geq (y_1 - ||y_*|| \cdot ||\bar{z}_*||)^2 t + ||y_*||^2 c\sqrt{t}.
\end{aligned}$$

Since y is an arbitrary unit vector, the Rayleigh-Ritz theorem implies that the minimum eigenvalue of $\sum_{s=1}^{t} \mathbb{E}_{s-1}[x_s x_s^\top]$ grows by least $\delta_0\sqrt{t}$ for some constant $\delta_0 > 0$. †

Claim 2. $|| \sum_{s=1}^{t} Var_{s-1}[x_s x_s^\top]|| \leq v_0\sqrt{t}$ *almost surely for some* $v_0 > 0$.

Proof of Claim 2

Let $C = \max_{i=1,\ldots,N}\{C_i\}$ and $R = Nu^2$ where $P = [l, u]$ is the admissible price interval with $0 < l < u$. By the definition of the conditional variance for random matrices:

$$
\begin{aligned}
\mathrm{Var}_{s-1}\left[x_s x_s^\top\right] &= \mathrm{Cov}_{s-1}\left[z\nu_s^\top, x_s x_s^\top\right] + \mathrm{Cov}_{s-1}\left[\nu_s z^\top, x_s x_s^\top\right] + \mathrm{Cov}_{s-1}\left[\nu_s\nu_s^\top, x_s x_s^\top\right] \\
&= \mathbb{E}_{s-1}\left[(\nu_s^\top \nu_s) z z^\top\right] + \mathbb{E}_{s-1}\left[(\nu_s^\top z) z \nu_s^\top\right] + \mathbb{E}_{s-1}\left[(\nu_s^\top \nu_s) z \nu_s^\top\right] \\
&\quad + \mathbb{E}_{s-1}\left[(\nu_s^\top z)\nu_s z^\top\right] + \mathbb{E}_{s-1}\left[(z^\top z)\nu_s\nu_s^\top\right] + \mathbb{E}_{s-1}\left[(\nu_s^\top z)\nu_s\nu_s^\top\right] \\
&\quad + \mathbb{E}_{s-1}\left[(\nu_s^\top \nu_s)\nu_s z^\top\right] + \mathbb{E}_{s-1}\left[(\nu_s^\top z)\nu_s\nu_s^\top\right] + \mathbb{E}_{s-1}\left[(\nu_s^\top \nu_s)\nu_s\nu_s^\top\right] \\
&\quad - \mathbb{E}_{s-1}\left[\nu_s\nu_s^\top\right]^2 \\
&= \mathbb{E}_{s-1}\left[||\nu_s||^2 x_s x_s^\top\right] + \mathbb{E}_{s-1}\left[(\nu_s^\top z) x_s x_s^\top\right] + \mathbb{E}_{s-1}\left[(\nu_s^\top z)\nu_s\nu_s^\top\right] \\
&\quad + (z^\top z)\mathbb{E}_{s-1}\left[\nu_s\nu_s^\top\right] - \mathbb{E}_{s-1}\left[\nu_s\nu_s^\top\right]^2.
\end{aligned}
$$

Apply the spectral norm to both sides of the above equation and use Jensen's inequality on each term to generate an upper bound,

$$
\begin{aligned}
||\mathrm{Var}_{s-1}\left[x_s x_s^\top\right]|| &\leq \mathbb{E}_{s-1}\left[||\nu_s||^2 ||x_s x_s^\top||\right] + \mathbb{E}_{s-1}\left[|\nu_s^\top z|\ ||x_s x_s^\top||\right] \\
&\quad + \mathbb{E}_{s-1}\left[|\nu_s^\top z|\ ||\nu_s\nu_s^\top||\right] + \ z^\top z\mathbb{E}_{s-1}\left[||\nu_s\nu_s^\top||\right] \\
&\leq R\left(\mathbb{E}_{s-1}\left[||\nu_s||^2\right] + 2\mathbb{E}_{s-1}\left[|\nu_s^\top z|\right] + ||\mathbb{E}_{s-1}\left[\nu_s\nu_s^\top\right]||\right) \\
&\leq R\left(\mathbb{E}_{s-1}\left[||\nu_s||^2\right] + 2|z|^\top\mathbb{E}_{s-1}\left[\nu_s\nu_s^\top\right]^{1/2} + ||\mathbb{E}_{s-1}\left[\nu_s\nu_s^\top\right]||\right) \\
&\leq \frac{N \cdot R \cdot C}{\sqrt{s}} + O(\frac{1}{\sqrt{s}}).
\end{aligned}
$$

The resulting matrix variance statistic is then

$$
\sum_{s=1}^{t} ||\mathrm{Var}_{s-1}[x_s x_s^\top]|| \leq \sum_{s=1}^{t} \frac{NR}{\sqrt{s}} + O(\frac{1}{\sqrt{s}}) \leq v_0\sqrt{t},
$$

for some $v_0 > 0$. †

Let $\delta = \frac{\delta_0}{2}\sqrt{t}$ and $v = v_0\sqrt{t}$ and apply the Matrix Freedman corollary,

$$\begin{aligned} N \cdot \exp\left(\frac{-\delta^2/2}{v^2 + R\delta/3}\right) &\geq \mathbb{P}\left\{\lambda_{\min}(\mathcal{J}_t) \leq -\delta + \sum_{s=1}^{t} \lambda_{\min}\left(\mathbb{E}_{s-1}[x_s x_s^\top]\right)\right\} \\ Ne^{-\kappa_1\sqrt{t}} &\geq \mathbb{P}\left\{\lambda_{\min}(\mathcal{J}_t) \leq \kappa_0\sqrt{t}\right\}. \end{aligned}$$

□

We now invoke the following result from Keskin and Zeevi [2014].

Lemma 2.4.4 (Keskin and Zeevi [2014], Lemma 3)**.** There exist finite positive constants ρ and k such that,

$$\mathbb{P}\left\{||\hat{\theta}_{it} - \theta_i|| > \delta, \lambda_{\min}(\mathcal{J}_t) \geq m\right\} \leq kt \exp\left(-\rho(\delta \wedge \delta^2)m\right),$$

for all $\delta, m > 0$ and all $t \geq 3$.

The proof approach in the result above combined with our earlier results then yield the following result for the overall estimation error.

Proposition 2.4.5. *If firms choose admissible strategies and there exist constants $c_i, C_i > 0$ such that $\frac{c_i}{\sqrt{t}} \leq \mathrm{Var}_{t-1}[\sigma_{it}] \leq \frac{C_i}{\sqrt{t}}$ almost surely for $i = 1, \ldots, N$ and $t < T$, then firm i's expected estimation error in period t is:*

$$\mathbb{E}_\theta^\sigma\left[||\theta_i - \vartheta_{it}||^2\right] = O\left(\frac{\log(t)}{\sqrt{t}}\right).$$

Proof. We break up the expectation into cases where the minimum eigenvalue of $\mathcal{J}_t$ is large with respect to the current time period and when it is small.

$$
\begin{aligned}
\mathbb{E}_\theta^\sigma\left[||\theta_i - \vartheta_{it}||^2\right] &= \int_0^\infty \mathbb{P}_\theta^\sigma\left(||\theta_i - \vartheta_{it}||^2 > x, \lambda_{\min}(\mathcal{J}_t) \geq \kappa_0\sqrt{t}\right) dx \\
&\quad + \int_0^\infty \mathbb{P}_\theta^\sigma\left(||\theta_i - \vartheta_{it}||^2 > x, \lambda_{\min}(\mathcal{J}_t) < \kappa_0\sqrt{t}\right) dx \\
&\leq \int_0^\infty \mathbb{P}_\theta^\sigma\left(||\theta_i - \hat{\theta}_{it}||^2 > x, \lambda_{\min}(\mathcal{J}_t) \geq \kappa_0\sqrt{t}\right) dx \\
&\quad + \int_0^{K_1} \mathbb{P}_\theta^\sigma\left(||\theta_i - \vartheta_{it}||^2 > x, \lambda_{\min}(\mathcal{J}_t) < \kappa_0\sqrt{t}\right) dx \\
&\leq \int_0^\infty \mathbb{P}_\theta^\sigma\left(||\theta_i - \hat{\theta}_{it}||^2 > x, \lambda_{\min}(\mathcal{J}_t) \geq \kappa_0\sqrt{t}\right) dx \\
&\quad + K_1 \mathbb{P}_\theta^\sigma\left(\lambda_{\min}(\mathcal{J}_t) < \kappa_0\sqrt{t}\right),
\end{aligned}
$$

where $K_1 = \max_{\theta,\theta' \in \Theta} ||\theta - \theta'||^2$. The first inequality is due to the fact that the estimation errors of the projected least squares estimate ϑ_{it} are bounded by K_1 and are weakly smaller than the estimation errors of $\hat{\theta}_{it}$.

In the proof of Keskin and Zeevi [2014], Theorem 2, the authors prove that Lemma 2.4.4 implies that

$$
\int_0^\infty \mathbb{P}_\theta^\sigma\left(||\theta_i - \hat{\theta}_{it}||^2 > x, \lambda_{\min}(t) \geq \kappa_0\sqrt{t}\right) dx = O(\frac{\log t}{\sqrt{t}}).
$$

Since the steps are identical to their analysis, they are omitted for brevity. Apply Lemma 2.4.3 to get the desired conclusion:

$$
\mathbb{E}_\theta^\sigma\left[||\theta_i - \vartheta_{it}||^2\right] \leq \frac{A_0 \log t}{\sqrt{t}} + \frac{A_1}{\sqrt{t}},
$$

for some constants $A_0, A_1 > 0$. Therefore, the expected estimation error in each period is order $O(\frac{\log t}{\sqrt{t}})$. □

2.5 Low Regret under Random Dithering Policies

Using the information result of the previous section, we can devise a best-response with random dithering strategy that achieves low regret due to learning.

Theorem 2.5.1. *If firms adopt best-response with random dithering strategies and there exist constants* $c_i, C_i > 0$ *such that* $\frac{c_i}{\sqrt{t}} \leq \mathrm{Var}_{t-1}[\sigma_{it}] \ \leq \frac{C_i}{\sqrt{t}}$ *almost surely for* $i = 1, \ldots, N$ *and* $t < T$*, then the worst-case regret due to learning is* $O(\sqrt{T}\log T)$ *for all firms.*

Proof. Consider the maximum t period contribution to regret due to learning for any firm. Let $i^* = \arg\max_{i \in I} \mathbb{E}_\theta^\sigma\left[(p_i^{NE} - p_{it})^2\right]$. Next, separate the regret due to learning into estimation error and forecast error and dithering as follows:

$$
\begin{aligned}
\mathbb{E}_\theta^\sigma\left[(p_{i^*}^{NE} - p_{i^*t})^2\right] &= \mathbb{E}_\theta^\sigma\left[(\varphi(\theta_{i^*}, p^{NE}) - \varphi(\vartheta_{i^*t}, p_t^e) - \nu_{i^*t})^2\right] \\
&\leq 2\mathbb{E}_\theta^\sigma\left[\left(\varphi(\theta_{i^*}, p^{NE}) - \varphi(\vartheta_{i^*t}, p^{NE})\right)^2\right] \\
&\quad + 2\mathbb{E}_\theta^\sigma\left[\left(\varphi(\vartheta_{i^*t}, p^{NE}) - \varphi(\vartheta_{i^*t}, p_t^e) - \nu_{i^*t}\right)^2\right] \\
&= 2\mathbb{E}_\theta^\sigma\left[\left(\varphi(\theta_{i^*}, p^{NE}) - \varphi(\vartheta_{i^*t}, p^{NE})\right)^2\right] \\
&\quad + 2\mathbb{E}_\theta^\sigma\left[\left(\varphi(\vartheta_{i^*t}, p^{NE}) - \varphi(\vartheta_{i^*t}, p_t^e)\right)^2\right] + 2\mathrm{Var}_\theta^\sigma(\nu_{i^*t}).
\end{aligned}
$$

The first term on the right hand side representing the estimation error component can be bounded using the mean value theorem

$$
\left|\left(\varphi(\theta_{i^*}, p^{NE}) - \varphi(\vartheta_{i^*t}, p^{NE})\right)\right| \leq \max_{\substack{\omega \in \Theta, j \in I, \\ k \in \{1,\ldots,N+1\}}} \sqrt{(N+1)}\left|\frac{\partial\varphi(\omega_j, p^{NE})}{\partial\omega_{jk}}\right| ||\theta_{i^*} - \vartheta_{i^*t}||.
$$

Therefore,

$$\begin{aligned}\mathbb{E}_\theta^\sigma\left[(p_{i^*}^{NE} - p_{i^*t})^2\right] &\leq 2K_0\mathbb{E}_\theta^\sigma\left[||\theta_{i^*} - \vartheta_{i^*t}||^2\right] \\ &\quad +2\mathbb{E}_\theta^\sigma\left[\left(\varphi(\vartheta_{i^*t}, p^{NE}) - \varphi(\vartheta_{i^*t}, p_t^e)\right)^2\right] + 2\,\mathrm{Var}(\nu_{i^*t}),\end{aligned}$$

where K_0 is defined by the maximization in the previous equation.

Let $\vartheta_{i^*t} = (a_{i^*t}, b_{i^*t}, c_{i^*jt})^\top$ and substitute the definition of $\varphi(\vartheta_{i^*t}, p_t^e)$ into the above equation.

$$\begin{aligned}\mathbb{E}_\theta^\sigma\left[(p_{i^*}^{NE} - p_{i^*t})^2\right] &\leq 2K_0\mathbb{E}_\theta^\sigma\left[||\theta_{i^*} - \theta_{i^*t}||^2\right] \\ &\quad +2\mathbb{E}_\theta^\sigma\left[\left(\sum_{j\neq i^*}\frac{c_{ijt}}{-2b_{it}}(p_j^{NE} - p_{jt}^e)\right)^2\right] + 2\,\mathrm{Var}(\nu_{i^*t}).\end{aligned}$$

Choosing the firm with the largest gap between the Nash equilibrium price and the forecasted price,

$$\begin{aligned}\mathbb{E}_\theta^\sigma\left[(p_{i^*}^{NE} - p_{i^*t})^2\right] &\leq 2\max_{\vartheta\in\Theta}\left\{\left(\frac{\sum_{j\neq i^*} c_{i^*jt}}{-2b_{i^*t}}\right)^2\right\}\max_{j\in I}\left\{\mathbb{E}_\theta^\sigma\left[\left(p_j^{NE} - p_{jt}^e\right)^2\right]\right\} \\ &\quad + 2K_0\mathbb{E}_\theta^\sigma\left[||\theta_{i^*} - \theta_{i^*t}||^2\right] + 2\,\mathrm{Var}(\nu_{i^*t}).\end{aligned}$$

Define $\Gamma := 2\max_{\theta\in\Theta}\left[\left(\frac{\sum_j \gamma_j}{-2\beta}\right)^2\right]$ and notice that $\Gamma < 1$ by both Proposition 2.2.1 and the stronger assumption that $|\sum_{j\neq i}\gamma_{ij}| \leq -\beta_i$ for all $i \in 1\ldots N$ and all $\theta \in \Theta$. Therefore, when firms use an admissible forecast scheme the regret due to learning is

$$\begin{aligned}\mathbb{E}_\theta^\sigma\left[(p_{i^*}^{NE} - p_{i^*t})^2\right] &\leq \Gamma\max_{j\in I}\left\{\mathbb{E}_\theta^\sigma\left[\left(p_j^{NE} - p_{jt}^e\right)^2\right]\right\} \\ &\quad + 2K_0\mathbb{E}_\theta^\sigma\left[||\theta_{i^*} - \theta_{i^*t}||^2\right] + 2\,\mathrm{Var}(\nu_{i^*t}).\end{aligned}$$

Applying Proposition 2.4.5, and the assumptions on the variance of the added noise,

$$\mathbb{E}^{\sigma}_{\theta}\left[(p^{NE}_{i^*} - p_{i^*t})^2\right] \;\leq\; \Gamma \max_{j\in I}\left\{\mathbb{E}^{\sigma}_{\theta}\left[\left(p^{NE}_{j} - p^{e}_{jt}\right)^2\right]\right\} + K\frac{\log t}{\sqrt{t}},$$

for some $K > 0$, determined by the bounds C_{i^*}, K_0, and the upper bound on the estimation error. Next, represent that the price forecasts in the above equation as a weighted sum of past prices,

$$\mathbb{E}^{\sigma}_{\theta}\left[(p^{NE}_{i^*} - p_{i^*t})^2\right] \;\leq\; \Gamma \max_{j\in I}\left\{\mathbb{E}^{\sigma}_{\theta}\left[\left(p^{NE}_{j} - \sum_{\tau=1}^{t-1}\lambda_\tau p_{j\tau}\right)^2\right]\right\} + K\frac{\log t}{\sqrt{t}}.$$

Applying Jensen's inequality to the sum,

$$\begin{aligned}\mathbb{E}^{\sigma}_{\theta}\left[(p^{NE}_{i^*} - p_{i^*t})^2\right] \;&\leq\; \Gamma \max_{j\in I}\left\{\mathbb{E}^{\sigma}_{\theta}\left[\sum_{\tau=1}^{t-1}\lambda_\tau\left(p^{NE}_{j} - p_{j\tau}\right)^2\right]\right\} + K\frac{\log t}{\sqrt{t}}\\ &\leq\; \Gamma\sum_{\tau=1}^{t-1}\lambda_\tau \max_{j\in I}\left\{\mathbb{E}^{\sigma}_{\theta}\left[\left(p^{NE}_{j} - p_{j\tau}\right)^2\right]\right\} + K\frac{\log t}{\sqrt{t}},\end{aligned}$$

which forms the autoregressive sequence

$$y_t \leq \Gamma\sum_{\tau=1}^{t-1}\lambda_\tau y_\tau + K\frac{\log t}{\sqrt{t}}.$$

with $y_t = \max_{i\in I}\mathbb{E}^{\sigma}_{\theta}\left[(p^{NE}_{i} - p_{it})^2\right]$. Note that the admissible forecast functions are Cournot adjustment, fixed H-horizon time averages and exponential decay. Each case generates an exponential decay in the lagged price terms, with decay rates varying according to the choices for the λ_τ terms. Summing across time, we attain our desired result through the integral bound.

$$\max_{i\in I}\mathbb{E}_\theta^\sigma\left[(p_i^{NE}-p_{it})^2\right] = O(\frac{\log t}{\sqrt{t}}) \quad \forall\theta\in\Theta$$
$$\Downarrow$$
$$\max_{\theta\in\Theta,i\in I}\sum_{t=1}^{T}\mathbb{E}_\theta^\sigma\left[(p_i^{NE}-p_{it})^2\right] = O(\sqrt{T}\log T).$$

□

2.6 Managerial Implications

These results imply that an equilibrium among firms using a policy of best-response against forecast plus noise can achieve complete learning and attain the full-information Nash equilibrium revenues if the noise terms satisfy the conditions above. In numerical results, however, if firms do not add sufficient noise then incomplete learning may occur and different effects can occur, including the possibility that all firms earn greater revenues than under the full-information Nash equilibrium. An example appears in Figure 2.1. The green and blue curves in that figure correspond to the regions that are favorable to each firm relative to the Nash equilibrium which occurs at the lower intersection of the two curves. The upper section corresponds to higher revenues for firm 1 and the right section corresponds to higher revenues for firm 2. The crosses correspond to the repeated actions of best responses with varying amounts of noise. The red and green trajectories correspond to sequences that terminate with a non-Nash equilibrium price with higher revenues for firm 1 than under the Nash equilibrium. The blue sequence of prices has little variation and terminates a point in which both firms have higher revenues than under the Nash equilibrium because their learning is incomplete.

This was also observed in the situation of unknown competition explored in Cooper et al. [2015]. This phenomenon suggests that it may be beneficial for firms not to experiment on

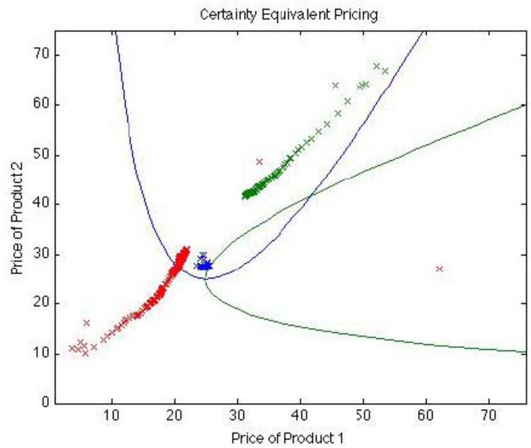

Figure 2.1: Incomplete learning price trajectories.

prices if the information from such revelation can lead to competitors' learning (perhaps free-riding) and reduced revenues. To make the analysis of such a situation tractable, the next chapter introduces a model with a simplified problem structure in the form of a repeated prisoner's dilemma game with a fully general policy structure and demonstrates that strategic lack of learning can occur in this setting, even with a finite time horizon.

CHAPTER 3

INCOMPLETE LEARNING AS AN EQUILIBRIUM STRATEGY

This chapter addresses the issue of incomplete learning under competition by directly investigating equilibrium pricing policies in dynamic games of incomplete information. The model shows that incomplete learning of demand information is not merely a byproduct of the dynamic pricing strategies investigated in the first model, but is rather a rational outcome for firms competing in a Markov perfect equilibrium (MPE). In particular, we develop several simple demand environments, where the equilibrium strategies for the competing firms actively avoid learning the true value of demand and attain a collusive outcome even in finite time horizons. Hence, firms prefer to remain willfully ignorant of the marketplace for their products.

By simplifying the information assumptions of Chapter 1, we present a dynamic game where the equilibria can be characterized using simple Markov strategies. Consider two firms that choose between two potential prices, the single period Nash equilibrium price and the cooperative equilibrium price. Initially, they do not know which price represents cooperation and which represents competition. In this respect, the competitive single period Nash equilibrium price and the cooperative equilibrium price form a prisoner's dilemma. Assuming a common Bayesian prior over the two possibilities, each firm decides whether or not to stick with the historically charged price or to experiment. Our results show that there exist conditions, parameterized by the value of the game's payoffs and prior beliefs, such that the firms will deliberately avoid experimenting with the prices in an effort to remain uncertain.

The significance of the results here are that firms can rationally choose not to experiment and learn the environment because their competitor also benefits from the information and can use that knowledge to the detriment of the experimenting agent. This threat of information leakage can dominate for any finite time horizon, leading to maintenance of a cooperative equilibrium and continued uncertainty over the environment. This contrasts with the traditional literature described below that opportunities for collusion are reduced in conditions with uncertainty (since deviations from cooperation are less likely to be detected and punished).

3.1 Literature Review

This model relates to the extensive literature on collusion and the mechanisms that sustain or disrupt it. As noted in the previous chapter, much of this literature extends from Stigler [1964] where lack of certainty is viewed as a classic deterrent to sustaining collusion. In later papers, Green and Porter [1984] and Abreu et al. [1990] showed that collusion is possible without full observability but the mechanism to sustain collusion is still based on some form of observed signals and punitive actions by competitors. The model here differs in that the threat to competitors is the information that can be revealed through exploratory action.

This model also builds on the extensive literature on prisoner's dilemma games, in which two competitors chose either to cooperate or defect. In this literature, cooperation is generally not sustainable with a finite horizon (see Benoit and Krishna [1985] and Marlats [2015]) or costly information (see Fong et al. [2008], Rahman [2014], Yamamoto [2015], and Yamamoto [2014]). These papers seek to develop Folk Theorems in which infinite horizon equilibria can sustain any set of feasible payoffs. Our result is different in that collusion arises in the finite horizon and where information can be obtained but remains hidden.

3.2 Model

Consider a stochastic game where two firms compete in a common market, with each selling a distinct product over a fixed T period selling season. The sequence of actions and outcomes in each period follow the same steps as our first model. Firms begin each period with a belief about an uncertain demand environment and privately choose a price for their product from a fixed set of feasible prices. The firms then announce their chosen prices publicly and simultaneously, and in turn are given private realizations of demand from their customers. The distinguishing feature of this model, compared to our previous OLS approach, is that the set of possible prices for each firm is limited to a finite set of actions, instead of a closed interval, and the parameters for the underlying demand are restricted to a set of discrete demand “scenarios”, rather than a compact set Θ.

Before this selling season begins, a state of the world is drawn from one of two possible market scenarios. For each state of the world, the profits to each firm are conditionally deterministic; that is, if a firm knew the true underlying market scenario, then that firm would know exactly what revenues would result from each possible pair of prices. In the first period, firms are unaware of the true market scenario and share a common prior π, which equals the probability that the underlying state is scenario 1. Specifically, each scenario corresponds to a different Prisoner's dilemma,

Scenario 1:

	a_1	a_2
a_1	X, X	T_1, S_1
a_2	S_1, T_1	R_1, R_1

Scenario 2:

	a_1	a_2
a_1	X, X	S_2, T_2
a_2	T_2, S_2	P_2, P_2

where $S_1 < X < R_1 < T_1$ and $S_2 < P_2 < X < T_2$.

At the start of each period, firms simultaneously choose between two prices: price a_1 and price a_2. When both firms elect to charge price a_1, the outcome is known to be a fixed value, X. This joint action reveals no new information and the game continues to the next period. However, if either of the firms chose to charge price a_2, then the true demand scenario is revealed to both firms after the revenues for that period are collected. For this reason, we refer to the pure-strategy of charging a_1 as *staying put* and the pure-strategy of charging a_2 as *experimenting*. The firms, therefore, have joint control of the information state of the game.

To solve for the Markov perfect equilibria of this game of uncertain information, we begin by considering the single period game, $T = 1$. Let p be the probability that firm 1, (the ROW player in the Prisoner's dilemma), chooses price a_1 and q be the probability that firm -1, (the COL player) chooses price a_2. Then for a given q the best response for firm 1 is

$$\mathcal{BR}(q) = \arg\max_{p\in[0,1]} \begin{pmatrix} pq \\ p(1-q) \\ (1-p)q \\ (1-p)(1-q) \end{pmatrix}^{\top} \left[\pi \begin{pmatrix} X \\ T_1 \\ S_1 \\ R_1 \end{pmatrix} + (1-\pi) \begin{pmatrix} X \\ S_2 \\ T_2 \\ P_2 \end{pmatrix} \right],$$

with the left vector representing the probability of each joint action and the vector sum on the right representing the expected value of each action for firm 1. Through the following transformations, the best-response function is fully represented by a 2-dimensional state of game coordinates $x_\pi \in \Re$ and $y_\pi \in \Re$.

$$\begin{aligned} \mathcal{BR}(q) = & \arg\max_{p\in[0,1]} \; p \begin{pmatrix} q \\ (1-q) \\ -q \\ -(1-q) \end{pmatrix}^{\top} \left[\pi \begin{pmatrix} X \\ T_1 \\ S_1 \\ R_1 \end{pmatrix} + (1-\pi) \begin{pmatrix} X \\ S_2 \\ T_2 \\ P_2 \end{pmatrix} \right] + \alpha(\pi, q) \\ = & \arg\max_{p\in[0,1]} p \begin{pmatrix} q \\ (1-q) \end{pmatrix}^{\top} \left[\pi \begin{pmatrix} X - S_1 \\ T_1 - R_1 \end{pmatrix} + (1-\pi) \begin{pmatrix} X - T_2 \\ S_2 - P_2 \end{pmatrix} \right] \\ = & \arg\max_{p\in[0,1]} p \begin{pmatrix} q \\ (1-q) \end{pmatrix}^{\top} \begin{pmatrix} x_\pi \\ y_\pi \end{pmatrix}, \end{aligned}$$

where

$$\begin{aligned} x_\pi &:= \pi(X - S_1) - (1-\pi)(T_2 - X) \\ y_\pi &:= \pi(T_1 - R_1) - (1-\pi)(P_2 - S_2) \\ \alpha(\pi, q) &:= q\,(\pi S_1 + (1-\pi)T_2) + (1-q)\,(\pi R_1 + (1-\pi)P_2)\,, \end{aligned}$$

with x_π representing the expected gain to firm 1 (given π) of price a_1 when the opponent chooses to stay put ($q = 1$) and y_π representing the expected gain of staying put when the opponent experiments ($q = 0$).

Written in a simplified form,

$$\mathcal{BR}(q) = \arg\max_p \alpha(\pi, q) + \beta(\pi, q)p \quad = \quad \begin{cases} 0, & \beta(q) \le 0, \\ 1, & \beta(q) > 0, \end{cases}$$

where $\beta(\pi, q)$ represents the expected payoff of charging a_1 and possibly remaining ignorant of demand, given the strategy of the opponent is q. Specifically, $\beta(\pi, q) = qx_\pi + (1-q)y_\pi$ is the q weighted convex combination of x_π and y_π. Firm 1's decision is then straightforward. If the expected payoff of charging a_1 is positive, then firm 1 charges a_1. If it is negative, then firm 1 experiments and chooses a_2. If it is zero, then firm 1 is indifferent between experimenting and staying put and could chose either action or adopt a mixed strategy. The following two dimensional graph illustrates the regions of the (x_π, y_π) graph and the pure-strategy equilibrium actions in each quadrant.

Each quadrant of the (x_π, y_π) represents a region where for a given belief π, the prescribed pure strategy actions are single-period Nash equilibria. For instance, if for a given π the values of x_π and y_π were both positive, then the joint action $(p^{NE}, q^{NE}) = (a_1, a_1)$ is a pure-strategy Nash equilibrium. Note that in quadrant II and quadrant IV of the graph above, mixed strategy equilibrium exist where both firms 1 and 2 are indifferent between

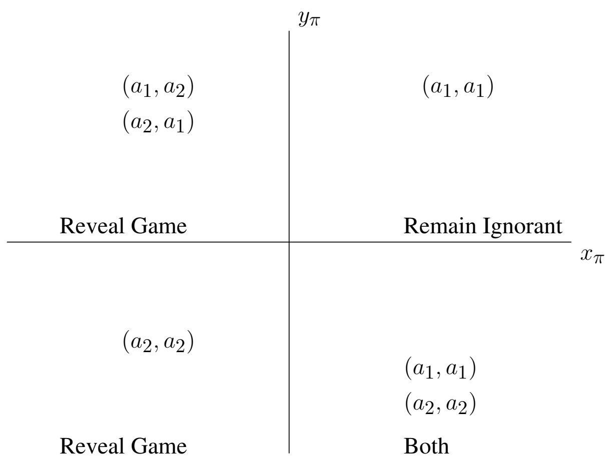

Figure 3.1: Pure-Strategy Nash EQ in Single-Period Game

learning the game and remaining uncertain.

Next, we extend the single-period stage game to a T-period dynamic game. We use a general setting of Markov perfect equilibrium as our solution concept. A Markov perfect equilibrium is a subgame perfect equilibrium where the firms are restricted to Markov strategies $\sigma(s) : S \to \mathcal{P}([0,1])^T$ where $S = \{0, \pi, 1\}$ denotes the set of three possible belief states for the firms. When $s = 1$, both firms are aware that scenario 1 represents the underlying payoff structure of the game and likewise, when $s = 0$, the game is known to be scenario 2. Let $\sigma_1(s) = (p_1(s), \ldots, p_T(s)) \in [0,1]^T$ and $\sigma_{-1} = (q_1(s), \ldots, q_T(s)) \in [0,1]^T$ denote the Markov strategies of firm 1 and -1, respectively, and the set Σ denote the set of all Markov strategies. The term p_t denotes the probability that firm 1 stays put in period t; q_t is analogously defined for firm -1. Given these strategies, the expected payoff to firm i for the subgame beginning in period t is denoted $u_i(\sigma, \sigma_{-i}, t)$. A pair of strategies form a Markov perfect equilibrium when

$$u_i(\sigma_i, \sigma_{-i}, t) \geq \max_{\bar{\sigma}_i \in \Sigma} u_i(\bar{\sigma}_i, \sigma_{-i}, t), \quad i \in \{1, -1\}, \forall t \in \{1, \ldots, T\}.$$

The space of Markov strategies Σ can be reduced significantly, as there are unique equilibria for the states $s = 1$ and $s = 0$. If either of these states are reached in period t, by either firm charging a_2 in period $t-1$, then the only subgame perfect equilibrium strategy is for both firms is to charge a_1 if $s = 1$ or a_2 if $s = 0$ for all periods t through T. This is due to the standard backwards induction argument for a finite horizon, prisoner's dilemma. Both firms know that their opponent has a dominant strategy, which is to choose these prices in the final period, and therefore have no incentive to cooperate in the period prior. However, if the game is uncertain in period t, we now show that the deterrent to price exploration enforces a more collusive outcome, due to the uncertainty between two prisoner's dilemma games.

The expected payoffs of each firm for the T-period games are determined using backwards induction. Let $V_\pi^\sigma[t]$ denote the expected payoff of firm i from periods t through T given that the game has not yet been revealed (i.e., $s = \pi$) by period t. Let $v_0 = \pi X + (1-\pi)P_2$ denote the single-period payoff that firm i expects, given his beliefs, the revealed game to yield in the next period. Then the expected payoff for firm i given strategies σ are

$$\begin{aligned} V_\pi^\sigma[t] &= SG_\pi(0) + \mathbb{E}[u_i(\sigma, t+1)] \\ &= SG_\pi(0) + p_t q_t V_\pi^\sigma[t+1] + (1 - p_t q_t) v_0 (T-t) \\ &= SG_\pi(0) + p_t q_t \left(V_\pi^\sigma[t+1] - v_0(T-t)\right) + v_0(T-t), \end{aligned}$$

where the value $SG_\pi(0)$ is equal to the expected payoff to firm i of the single-period stage game that was analyzed previously. Using the (x_π, y_π) graphical analysis to characterize equilibria strategies, let $SG(\Delta)$ represent the expected payoff to firm i of an equilibrium strategy to an augmented single-period game with $(x', y') = (x_\pi + \Delta, y_\pi)$ as the game coordinates. That is, plot the point (x', y') on the graph in figure 3.1 and select a pure-strategy Nash equilibrium (p^{NE}, q^{NE}) from the appropriate quadrant. The value $SG(\Delta)$ is then the expected value of this augmented game for firm i.

$$\begin{aligned} SG_\pi(\Delta) &:= \left\{u(p^{NE}, q^{NE}, T) : (x', y') = (x_\pi + \Delta, y_\pi)\right\} \\ &= \alpha(\pi, q^{NE}) + \left(q^{NE}x' + (1-q^{NE})y'\right)p^{NE} \\ &= \alpha(\pi, q^{NE}) + \left(q^{NE}(x_\pi + \Delta) + (1-q^{NE})y_\pi\right)p^{NE}. \end{aligned}$$

Since there exist multiple equilibria for certain quadrants in the (x_π, y_π) graph, the value of $SG_\pi(\Delta)$ is not well-defined without an explicit equilibria selection criterion. However, this formulation allows us to identify the Markov perfect equilibria to the dynamic game as follows. Consider the expected payoff equation for firm i introduced earlier:

$$\begin{aligned} V_\pi^\sigma[t] &= SG_\pi(0) + p_t q_t \left(V_\pi^\sigma[t+1] - v_0(T-t)\right) + v_0(T-t) \\ &= SG_\pi(0) + p_t q_t \Delta_{t+1} + v_0(T-t) \\ &= SG_\pi(\Delta_{t+1}) + v_0(T-t). \end{aligned}$$

The term $\Delta_{t+1} := V_\pi^\sigma[t+1] - v_0(T-t)$ represents the expected future gains when both firms stay put in period t less the gains from playing a complete information game for the next $T-t$ periods. Hence, the multi-period game is solved by reducing the decision in each period t for each firm to a single-period game. The future value of either knowing the scenario information or remaining uncertain is incorporated appropriately into this augmented game. Using this technique, we identify the following payoff and belief conditions where the strategy of $(p_t, q_t) = (1, 1)$ form a Markov perfect equilibrium, and both firms choose to stay put for the entire $T-$period game. Inefficient learning is then an equilibrium outcome for competing firms under these conditions.

Consider the case where $x_\pi > 0$; that is, the expected payoff of charging a_1 is positive when firm i knows that their opponent will also charge a_1 in a single-period game. In this case, firms would choose to remain ignorant of the underlying game in the terminal period T,

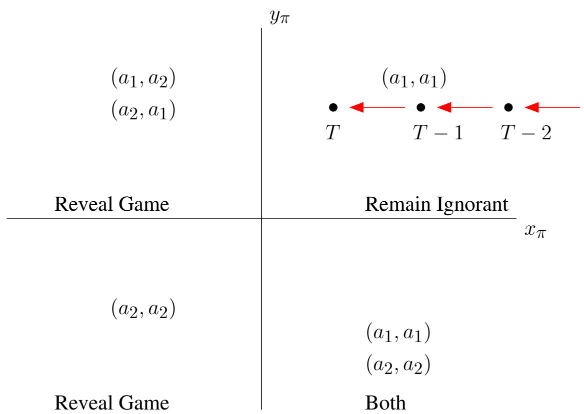

Figure 3.2: Inefficient Learning MPE in the Dynamic Game

provided that firms reached that period without changing their prices. Iterating backwards, we show that the added value of remaining ignorant persists in each period and as a result $\Delta_t > \Delta_{t+1}$ for all t and firms choose (a_1, a_1) in all periods. This process is illustrated in figure 3.2. Essentially, the analysis shows that if firms are willing to remain ignorant in the terminal period, they will also be willing to remain ignorant over an arbitrarily long, finite horizon.

Proposition 3.2.1 (Equilibrium with Incomplete Learning)**.** *If the set of payoff parameters and beliefs are such that* $x_\pi > 0$*, then the strategy* $(p_t, q_t) = (1, 1)$ *for all* $t \in \{1, \ldots, T\}$ *is a Markov perfect equilibrium.*

Proof. First recall from the analysis in figure 3.1 that when $x_\pi > 0$ the subgame beginning in period T has a Nash equilibrium $(p_T, q_T) = (1, 1)$. Hence, if the firms were to choose the action (a_1, a_1) for all periods $t < T$ and remain ignorant of the underlying game payoffs, they would continue to choose (a_1, a_1) in the terminal period. Continuing with this graphical approach, we next show that $\Delta_t \geq 0$ for $t < T$ and thus $x' \geq x_\pi$ for these periods as well. The backwards induction argument proceeds as follows. First, consider the terminal case where $\Delta_T = \alpha(\pi, 1) + x_\pi - v_0$. Substituting the definitions of α and x_π,

$$\begin{aligned} \Delta_T &= \alpha(\pi, 1) + x_\pi - v_0 \\ &= (1 - \pi)(X - P_2) > 0. \end{aligned}$$

Next, assume for a given $t < T$ that $(p_\tau, q_\tau) = (1, 1)$ and $\Delta_\tau \geq 0$ for $\tau > t$. Since $\Delta_{t+1} \geq 0$ the game coordinates $(x_\pi + \Delta_{t+1}, y_\pi)$ lie in quadrant I or quadrant IV of the single-period equilibrium graph. Therefore, the action $(p_t, q_t) = (1, 1)$ is an equilibrium action for period t. It remains to show that $\Delta_t \geq 0$.

$$\begin{aligned} \Delta_t &= V_\pi^\sigma[t] - v_0(T - t + 1) \\ &= (\alpha(\pi, 1) + x_\pi)(T - t + 1) - v_0(T - t + 1) \\ &= \Delta_T(T - t + 1) > 0. \end{aligned}$$

□

3.3 Competition with Demand and Opponent Uncertainty

In the model of the previous section, firms receive the same private information in each round and are effectively aware of their competitor's information state. In this section, we introduce uncertainty into the payoff signals and, hence, distinct private information states. We show that, although firms can become aware of the game environment, they may still have an incentive to sustain the cooperative equilibrium.

Consider two games with the following payoff structure:

Scenario 1:

	a_1	a_2
a_1	X_1^t, X_{-1}^t	T_1, S_1
a_2	S_1, T_1	R_1, R_1

Scenario 2:

	a_1	a_2
a_1	X_1^t, X_{-1}^t	S_2, T_2
a_2	T_2, S_2	P_2, P_2

where

$$\begin{array}{ccccccccc} S_2 & < & P_2 & < & \left[\begin{array}{c} \operatorname{supp} X_1^t \cup X_{-1}^t \\ \operatorname{supp} X_1^t \cup X_{-1}^t \end{array} \right] & < & T_2 & & \\ & & S_1 & < & & < & R_1 & < & T_1, \end{array}$$

for all periods $t \leq T$. The key distinction between these scenarios and the game considered in the previous section is that the rewards for playing action (a_1, a_1) are independent Bernoulli random variables. If the scenario is equal to 1 then,

$$X_i^t \sim \begin{cases} X_0, & \text{w.p. } \gamma \\ X_0 \pm \delta_1, & \text{w.p. } \frac{1}{2}(1-\gamma_t) \end{cases} \qquad i = \pm 1, t = 1, \ldots, T.$$

If the scenario is equal to 2 then,

$$X_i^t \sim \begin{cases} X_0, & \text{w.p. } \gamma \\ X_0 \pm \delta_2, & \text{w.p. } \frac{1}{2}(1-\gamma_t) \end{cases} \qquad i = \pm 1, t = 1, \ldots, T,$$

for $\delta_1 \neq \delta_2$. There is a $(1-\gamma_t)$ probability in each period that firm i will learn the underlying game scenario, even if both firms choose the typically uninformative action of pricing (a_1, a_1). Furthermore, firm i may learn the game, with a realization of $X_0 \pm \delta_1$ or $X_0 \pm \delta_2$, while its opponent realizes X_0 and remains ignorant. These scenarios generate a partially observable stochastic game, where firms face uncertainty over the game outcomes and uncertainty about their opponent's information state. Each firm i maintains a private state $s_{it} \in \{1, \ldots, 5\}$ for

$t \in \{1, \ldots, T\}$. These states are defined as follows:

$$s_{it} = \begin{cases} 1 & \text{Game Revealed - Scenario 1} \\ 2 & \text{Game Revealed - Scenario 2} \\ 3 & \text{Game Not Revealed - all } X_0 \\ 4 & \text{Game Not Revealed - occurrence of } X_0 \pm \delta_1 \\ 5 & \text{Game Not Revealed - occurrence of } X_0 \pm \delta_2. \end{cases}$$

States 1 and 2 indicate that at some time $\tau < t$, one of the firms charged a_2 and thus both firms are aware of the underlying game scenario and are aware that their opponent also has this knowledge. The remaining states all correspond to the setting where both firms charged price a_1 for the past $t-1$ periods. Therefore, neither firm knows whether their opponent has seen a revealing outcome, $(X_0 \pm \delta_1)$ or $(X_0 \pm \delta_2)$, in any of the past periods. State 3 corresponds to the situation where firm i is unaware of the game scenario, and states 4 and 5 correspond to firm i's knowledge of scenario 1 or scenario 2 respectively.

Denote the strategy of firm i as $\sigma_{it}(s_i)$. To solve for a Markov perfect equilibrium, we assume that firm i knows its current state in each period and best-responds to a known strategy matrix $\sigma_{-it}(s)$ for $s \in \{1, \ldots, 5\}$ of its opponent. The main result of this section is that for a range of values for the common prior $\pi_0 = Pr(\text{Scenario } 1 | s_{i1} = 3)$, the horizon length T, and the probability of uninformative outcomes γ_t, there exists Markov perfect equilibria where firms choose (a_1, a_1) in all periods. In other words, the theorem states that cooperation persists when the game has not been revealed, even when one firm or both firms has discovered the environment. Each firm's uncertainty over the private state of their opponent is enough to enforce cooperation throughout the game.

First, we update some of the notation used in the previous model to include the addition of opponent uncertainty. Let s_{1t} as the vector $p_t \in [0,1]^5$ with each component s representing the probability that firm 1 will choose price a_1 in period t and state s. Likewise, represent s_{-1t} as the vector $q_t \in [0,1]^5$. There are two belief sequences for firm 1. The belief matrix, $\mu^t_{s\bar{s}} = Pr(s_{-1t} = \bar{s}|s_{1t} = s)$, maintains firm 1's belief in period t of its opponents state. Additionally, let $\pi_t = Pr(\text{Scenario } 1|s_{it} = 3)$ denote firm 1's belief that the underlying environment is scenario 1, provided this information is still unknown. Applying the same game coordinate system that was used in the previous model, define

$$\begin{aligned} x_{\pi_t} &:= \pi_t(X_0 - S_1) - (1-\pi_t)(T_2 - X_0) \\ y_{\pi_t} &:= \pi_t(T_1 - R_1) - (1-\pi_t)(P_2 - S_2). \end{aligned}$$

Incorporating these new factors into the analysis used in the previous section yields the following theorem.

Theorem 3.3.1. *There exists a Markov perfect equilibrium strategy with* $p_t = q_t = (1,0,1,1,1)^\top$ *for* $t \in \{1,\dots,T-1\}$, *and* $p_T = q_T = (1,0,1,1,0)^\top$ *provided that*

i. $x_{\pi_0} > 0$;

ii. $\prod_{t=1}^{T-1} \gamma_t > \left(\frac{T_2 - X_0}{T_2 - P_2}\right)$.

Proof. Note that the only equilibrium action in state 1 is for both firms to choose (a_1, a_1) and similarly firms choose (a_2, a_2) in state 2. Each of these states generate finite horizon, complete-information games with dominant Nash equilibrium strategies. As a result, $p_{1t} = q_{1t} = 1$ and $p_{2t} = q_{2t} = 0$ are the only potential equilibrium actions for these revealed states.

The action (a_1, a_1) is also a dominant strategy for state 4. To see this, assume that firm 1 is in state 4 at period t and therefore knows that the underlying game is scenario 1, but is unaware of its opponent's state. Given knowledge of the underlying game, firm 1 has no myopic incentive to charge a_2 since $S_1 < R_1 < x_0$. Firm 2 also has no information-state incentive to charge a_2 as the value of the realized game (state 1) is the same as the value of obscured game (state 4) provided $q_{1t} = q_{3t} = q_{4t} = 1$. Firm 1's best response is therefore to charge a_1 for the remaining horizon.

Assuming that firm -1 adopts q_t as its strategy, the best-response strategy for firm 1 while in state 3 and 5 is determined using the analysis of the previous section. In the terminal period T, the best-response for firm 1 in state 3 is

$$\begin{aligned}\mathcal{BR}_3(q_T) &= \arg\max_p \alpha(\pi_T, \mu_3.q_T) + \beta(\pi_T, \mu_3.q_T)p \\ &= \arg\max_p \left[(\mu_{33}^T + \mu_{34}^T)x_{\pi_T}\right] p.\end{aligned}$$

By assumption $x_{\pi_0} > 0$ and $q_{3t} = q_{4t} = q_{5t} = 1$ for $t < T$. Applying Bayes rule to the prior belief on scenarios, $\pi_t = Pr(Scenario = 1|s_{1t} = 3)$ yields

$$\pi_{t+1} = \pi_t \frac{\mu_4^T.q_t}{\mu_3^T.q_t} = \pi_t,$$

and thus $\pi_T = \pi_0$. As a result, $(\mu_{33}^T + \mu_{34}^T)x_{\pi_0} > 0$ and the best-response for firm 1 in state 3 and period T is to charge the uninformative price a_1. If firm 1 were instead in state 5 in the terminal period T, the best-response to q_T is

$$\begin{aligned}\mathcal{BR}_5(q_T) &= \arg\max_p \alpha(0, \mu_5.q_T) + \beta(0, \mu_5.q_T)p \\ &= \arg\max_p \left[(\mu_{53}^T + \mu_{54}^T)x_0\right] p.\end{aligned}$$

By definition, $x_0 := -(T_2 - X_0) < 0$, and therefore firm 1 would choose price a_2 in the terminal period. To simplify the notation for the backwards induction, let

$$\begin{aligned}v_0 &:= \pi_0 V_{1,T} + (1 - \pi_0)V_{2,T} \\ \Delta_{3,t+1} &:= V_{3,t+1} - v_0 \\ \Delta_{5,t+1} &:= V_{5,t+1} - V_{2,t+1}.\end{aligned}$$

The terms $V_{s,t}$ represent the expected payoffs to firm 1 of the subgame starting in period t with information state s, provided that firm -1 chooses strategy q_t. Note that $\Delta_{3,T} = (1-\pi_0)(X_0 - P_2) > 0$ is the single-period expected gain from being ignorant of both the underlying game and the opponent's information-state, and $\Delta_{5,T} = \mu_{53}^T(T_2 - P_2)$ is the single-period expected gain from being ignorant of the opponent's information-state.

Iterating backwards, assume that both $\Delta_{3,t+1}$ and $\Delta_{5,t+1}$ are positive and consider the value function for state 3,

$$\begin{aligned} V_{3,t} &= \max_p \alpha(\pi_0, 1) + \beta(\pi_0, 1)p \\ &\quad + p\left(\gamma_t V_{3,t+1} + (1-\gamma_t)\left[\pi_0 V_{4,t+1} + (1-\pi_0)V_{5,t+1}\right]\right) \\ &\quad + (1-p)\left(\pi_0 V_{1,t+1} + (1-\pi_0)V_{2,t+1}\right). \end{aligned}$$

In words, the value of state 3 in period t is equal to the sum of myopic stage-game value $\alpha(\pi_0, 1) + \beta(\pi_0, 1)p$, the expected value of the subgame in period $t+1$ if firm 1 were to charge a_1, and the expected value of the subgame if firm 1 were to charge a_2. Simplifying the above equation,

$$V_{3,t} = \max_p \left(x_{\pi_0} + \gamma_t \Delta_{3,t+1} + (1-\gamma_t)(1-\pi_0)\Delta_{5,t+1}\right)p + \alpha(\pi_0, 1) + v_0(T-t).$$

From condition (1) of the theorem, $x_{\pi_0} > 0$ and by the inductive hypothesis $\Delta_{3,t+1} > 0$ and $\Delta_{5,t+1} > 0$. Therefore, the strategy $p_{3,t} = 1$ is the best response to q_t. Next, consider the value function for state 5.

$$\begin{aligned} V_{5,t} &= \max_p \alpha(0,1) + \beta(0,1)p + pV_{5,t+1} + (1-p)V_{2,t+1} \\ &= \max_p \left(\beta(0,1) + \Delta_{5,t+1}\right)p + \alpha(0,1) + V_{2t+1} \\ &= \max_p \left(x_0 + \Delta_{5,t+1}\right)p + \alpha(0,1) + V_{2t+1}. \end{aligned}$$

By the analysis above, $p_{5,t} = 1$ is the best response strategy to q_t provided that $\Delta_{5,t+1} > -x_0$. Note that this condition holds in the terminal period by condition (2) of the theorem.

$$\begin{aligned}
\Delta_{5,T} &= \mu_{53}^{T}(T_2 - P_2) \\
&= (T_2 - P_2)\prod_{t=1}^{T-1}\gamma_t \\
&> T_2 - x_0 \\
&> -x_0.
\end{aligned}$$

It remains to show that $\Delta_{3,t} > 0$ and that $\Delta_{5,t} > -x_0$. Assume that this condition holds for $\tau \geq t$. Then,

$$\begin{aligned}
\Delta_{5,t} &= V_{5,t} - V_{2,t} \\
&= \Delta_{5,t+1} + x_0 + \alpha(0,1) - P_2 \\
&= \Delta_{5,t+1} + (X_0 - T_2) + T_2 - P_2 \\
&= \Delta_{5,t+1} + (X_0 - P_2) \\
&> \Delta_{5,t+1} \\
&> -x_0,
\end{aligned}$$

and

$$\begin{aligned}
\Delta_{3,t} &= V_{3,t} - (\pi_0 V_{1,t} + (1-\pi_0)V_{2,t}) \\
&= \gamma_t \Delta_{3,t+1} + (1-\gamma_t)(1-\pi_0)\Delta_{5,t+1} + x_{\pi_0} + \alpha(\pi_0, 1) - v_0 \\
&> x_{\pi_0} + \alpha(\pi_0, 1) - v_0 \\
&= X_0 - v_0 \\
&= X_0 - (\pi_0 X_0 + (1-\pi_0)P_2) > 0.
\end{aligned}$$

□

3.4 Conclusion

This chapter and the last presented dynamic models of price competition with unknown demand. In the previous chapter, we presented a model in which the firms assume linear demand functions and can observe others' actions but not their payoffs. We showed that, in the class of strategies characterized by best responses plus noise, an outcome exists in which each firm adjusts their noise to be sufficient for learning and decreasing to ensure convergence. We also gave an example to show that it is possible, however, for other strategies to dominate this outcome for all firms. The lack of information, however, makes it difficult for the firms to coordinate on this outcome.

The second model in this chapter considered relaxation of the policies to obtain an equilibrium in which learning does not occur and firms jointly attain a cooperative outcome. To make this analyzable, the second model restricts the uncertainty to finite unknown states of the world and actions in a dynamic prisoner's dilemma situation. When firms' actions can reveal the state to their competitor and the payoffs overlap in a way that does not reveal the state with other actions, an equilibrium can result in which the firms maintain non-revealing actions and achieve the cooperative outcome for any time horizon. This result can even occur after an agent becomes informed.

These results imply that firms in competition can achieve learning as efficiently as a monopolist firm, but this requires either restriction in strategies or information structure. With sufficiently diffuse information but rapid learning by all participants from variations in actions, equilibria can exist in which learning does not occur. This result suggests that limitations on actions (such as laws, particularly on insurance products, that restrict price adjustments or price caps) and privacy restrictions can lead to inefficient outcomes (or to losses in consumer welfare).

This work suggests several potential follow-on studies. Further theoretical studies could consider general conditions for the non-learning phenomenon (or for its non-existence) in a broader game context. Mechanisms to reduce the possibilities of inefficient outcomes, such as the presence of a monitor or intermediary, could also be considered. For empirical extensions, as mentioned pharmaceutical examples could be studied for the observed phenomenon. In addition, the effects of price restrictions as in insurance markets might also provide useful empirical research.

CHAPTER 4

THE SLATER CONUNDRUM

4.1 Introduction

Duality is one of the most useful tools for modeling and solving optimization problems. Properties of the dual problem are used to characterize the structure of optimal solutions and design algorithms. The dual may be easier to solve than the primal and there exist well-known sufficient conditions, such as Slater's constraint qualification, that imply *zero duality gap* between the optimal values of the primal and dual.

Many real world applications are naturally modeled in infinite dimensions, with examples in revenue management (Gallego and Van Ryzin [1994]), procurement (Manelli and Vincent [1995]) inventory management (Adelman and Klabjan [2005]), territory division and transportation (Carlsson [2012]).

In each case, the authors apply properties of the dual to simplify the problem and develop structure for optimal solutions.

The focus of our research is on the connection in infinite dimensional convex optimization between constraint qualifications for zero duality gap and the existence of optimal dual solutions that are easily characterized and have a meaningful economic interpretation. A precise overview of our main results is found in Section 4.1.2.

To motivate our results, we begin with concrete examples. These examples illustrate how the intuitions and interpretations common in finite dimensions do not necessarily hold in infinite dimensions. Duality concepts that do extend to infinite dimensional problems are more subtle and difficult to apply in practice. This development is meant to be accessible to readers with little or no background in functional analytic approaches to convex optimization.

4.1.1 Motivation

In finite-dimensional convex optimization, conditions for zero duality gap (such as Slater's constraint qualification) are well understood. Moreover, there is a standard form and economic interpretation of the dual. Researchers in operations research and economics typically define a vector of dual prices that index the (finitely many) constraints. Each price is interpreted as the marginal value of the constraining resource of the corresponding constraint. A real vector of dual prices in finite dimensional convex optimization is a convenient representation of a linear functional defined over the constraint space, termed a *dual functional*. It is this notion of dual functionals, possibly no longer representable as a real vector, that extends to infinite dimensional problems and allows us to build a duality theory for convex optimization problems over arbitrary ordered vector spaces. In the economics literature these dual functionals are called *prices* (see, for instance Debreu [1954]). This is valid terminology since dual functionals "price" constraints by mapping vectors in the constraint space to the real numbers. When there is a zero duality gap between an optimization problem and its Lagrangian dual, this pricing interpretation gives significant insight into the structure of the primal optimization problem. However, in infinite dimensions the connection between conditions for zero duality gap and interpretations of the dual are more complex. The following examples illustrate this complexity.

Example 4.1.1.

Consider the finite dimensional linear program

$$
\begin{aligned}
&\min x_1 \\
&\qquad x_1 \geq -1, \\
&\quad -x_2 \geq 0, \\
&x_1 - \frac{1}{i}x_2 \geq 0, \quad i = 3, 4, \ldots, 10.
\end{aligned}
\tag{4.1.1}
$$

The vector spaces used in this problem are easily characterized. The *primal variable space* $X = \mathbb{R}^2$ contains the feasible region and the *constraint space* $Y = \mathbb{R}^{10}$ contains the problem data for the primal constraints and is ordered by the cone

$$\mathbb{R}^{10}_+ = \left\{ y \in \mathbb{R}^{10} : y_i \geq 0,\, i = 1, \ldots, 10 \right\}.$$

The dual feasible region is a subset of the vector space Y' of linear functionals that map Y into $\mathbb{R}$ (see Luenberger [1969]). The space Y' is called the *algebraic dual* of Y and elements of Y' are called *dual functionals*. The dual constraints of (4.1.1) are

$$\psi((1, 0, 1, \ldots, 1)) = 1, \tag{4.1.2}$$

$$\psi((0, -1, -1/3, \ldots, -1/10)) = 0, \tag{4.1.3}$$

$$\psi \in (\mathbb{R}^{10})'_+. \tag{4.1.4}$$

Consider the dual constraint (4.1.2) and let $e_i \in \mathbb{R}^{10}$ be the vector that equals 1 in the i^{th} component and 0 otherwise. Since every dual functional ψ is linear, we write (4.1.2) as

$$
\begin{aligned}
\psi((1, 0, 1, \ldots, 1)) &= \psi(1 \cdot e_1 + 0 \cdot e_2 + \cdots + 1 \cdot e_{10}) \\
&= \psi(1 \cdot e_1) + \psi(0 \cdot e_2) + \cdots + \psi(1 \cdot e_{10}) \\
&= 1 \cdot \psi(e_1) + 0 \cdot \psi(e_2) + \cdots + 1 \cdot \psi(e_{10}).
\end{aligned}
$$

Therefore, ψ can be represented by a real vector $\psi_i := \psi(e_i)$ for $i = 1, \ldots, 10$. Using this notation, the dual of (4.1.1) is

$$\begin{aligned} \max -\psi_1 & \\ \psi_1 + \psi_3 + \psi_4 + \cdots + \psi_{10} =& 1, \\ -\psi_2 - (\psi_3/3) - (\psi_4/4) - \cdots - (\psi_{10}/10) =& 0, \\ \psi_i \geq & 0 \quad i = 1, 2, \ldots 10. \end{aligned} \tag{4.1.5}$$

Representing the dual functional ψ as the real vector $(\psi_1, \ldots, \psi_{10})$ is standard practice in finite dimensional optimization. Problem (4.1.5) is a finite dimensional linear program with a simple structure: the constraint matrix of the dual program is the *transpose* of the constraint matrix of the primal. This property may not hold in infinite dimensional problems, as seen in Example 4.1.2 below.

If the primal is feasible and bounded, then there is an optimal primal and dual solution with zero duality gap. An optimal primal solution is $(x_1^*, x_2^*) = (-1, -10)$ and an optimal dual solution is $(\psi_1^*, \ldots, \psi_{10}^*) = (1, 0, \ldots, 0)$, each with an objective value of -1.

The representation of the dual functional ψ as the real vector $(\psi_1, \ldots, \psi_{10})$ is convenient for interpreting the dual. The optimal value of ψ_i is the increase in the primal objective value resulting from one unit increase in the right-hand side of the i^{th} primal constraint. Therefore, when the value of the primal and dual are equal, optimal dual prices are the marginal prices a decision maker is willing to pay to relax each primal constraint. This relationship between dual functionals and the pricing of constraints is a key modeling feature behind the success of duality theory of finite dimensional optimization. $\triangleleft$

The next example demonstrates that many of the nice properties of finite dimensional optimization in Example 4.1.1 may fail to hold in infinite dimensions.

Example 4.1.2 (Karney [1981], Example 1)**.**

Consider the extension of Example 4.1.1 to infinitely many constraints but still finitely many variables.

$$
\begin{aligned}
\inf x_1 & \\
x_1 &\geq -1, \\
-x_2 &\geq 0, \\
x_1 - \frac{1}{i} x_2 &\geq 0, \quad i = 3, 4, \ldots.
\end{aligned}
\tag{4.1.6}
$$

The left-hand-side column vectors $(1, 0, 1, 1, \ldots)$ and $(0, -1, -1/3, -1/4, \ldots)$ and right-hand-side vector $(-1, 0, \ldots)$ belong to many choices of constraint space. This is typical in infinite dimensions where multiple nonisomorphic vector spaces are consistent with the constraint data. By contrast, in finite dimensions all finite dimensional vector spaces of dimension m are isomorphic to $\mathbb{R}^m$.

A natural choice for the constraint space of (4.1.6) is the vector space $\mathbb{R}^{\mathbb{N}}$ of all real sequences. A naive approach to formulating the dual of (4.1.6) when the constraint space is $\mathbb{R}^{\mathbb{N}}$ is to mimic the logic of Example 4.1.1: assign a real number ψ_i to each constraint $i = 1, 2, \ldots$, and then take the transpose of the primal constraint matrix. Assuming this representation is valid, define a dual functional ψ on the constraint space of (4.1.6) by an infinite sequence $\{\psi_i\}_{i=1}^{\infty}$ where $\psi_i := \psi(e_i)$ with e_i, again having 1 in the i^{th} component and 0 elsewhere. However, this representation of an arbitrary dual functional ψ is not valid unless the following condition holds. Let $\{a_i\}_{i=1}^{\infty}$ be an arbitrary vector in the constraint space. We say a dual functional ψ is *countably additive* if $\psi\left(\{a_i\}_{i=1}^{\infty}\right) = \sum_{i=1}^{\infty} a_i \psi_i$. Fortunately, when the constraint space vector space is $\mathbb{R}^{\mathbb{N}}$, all positive dual functionals are countably additive. Indeed, Basu et al. [2014] prove that positive dual functionals in the algebraic dual of $\mathbb{R}^{\mathbb{N}}$ can be expressed as positive sequences with finite support; that is, $\psi_i > 0$ for only finitely many $i \in \mathbb{N}$. Clearly, such dual functionals are countably additive. The dual program derived by taking the transpose of the constraint matrix in (4.1.6) is

$$
\begin{aligned}
\sup -\psi_1 & \\
\psi_1 + \psi_3 + \psi_4 + \cdots &= 1, \\
-\psi_2 - (\psi_3/3) - (\psi_4/4) - \cdots &= 0, \\
\psi_i &\geq 0, \quad i = 1, 2, \ldots \\
\psi_i &> 0, \quad \text{for at most finitely many } i.
\end{aligned}
\tag{4.1.7}
$$

Notice that (4.1.7) closely resembles its finite dimensional analogue (4.1.5). The justification for this dual formulation is based on countable additivity and is formalized in Section 4.3.1 (see (4.3.5)). Unlike the finite case, a duality gap exists between (4.1.6) and (4.1.7). The optimal solution to the primal program (4.1.6) is $(x_1^*, x_2^*) = (0, 0)$ with a value of 0. Since every feasible solution to the dual program (4.1.7) is nonnegative, the second constraint implies ψ_i is zero for all $i \in \{2, 3, \ldots\}$. Therefore, the optimal dual solution is $(1, 0, 0, \ldots)$ with a value of -1.

The existence of a duality gap may be surprising to readers familiar with Slater's constraint qualification in finite dimensions. Slater's result states that if there exists a primal solution x such that each constraint is strictly satisfied, then there is a zero duality gap. There are feasible points of (4.1.6) ($\bar{x} = (1, -1)$ for instance) that strictly satisfy each constraint and the existence of a duality gap appears to contradict Slater's constraint qualification. How can this be? In infinite dimensions, feasible points that strictly satisfy each constraint are not necessarily interior points of the positive cone of the constraint space, the condition required by infinite dimensional versions of Slater's result (see Theorem 4.2.1). Indeed, the positive cone of the vector space $\mathbb{R}^{\mathbb{N}}$ has an empty interior under every linear topology (Aliprantis and Border [2006]).

Alternatively, choose the constraint space of (4.1.6) to be the vector space ℓ_∞ of bounded sequences. This is a valid constraint space because the columns and the right-hand side of

(4.1.6) lie in ℓ_∞. The resulting dual program is

$$\begin{aligned} \sup \psi(\{-1,0,\dots\}) \\ \psi(\{1,0,1,1,\dots\}) = 1, \\ \psi(\{0,-1,-\frac{1}{3},-\frac{1}{4},\dots\}) = 0, \\ \psi \in (\ell_\infty)'_+. \end{aligned} \tag{4.1.8}$$

When the vector space ℓ_∞ is equipped with its norm topology, its positive cone has a non-empty interior. The infinite dimensional version of Slater's constraint qualification (Theorem 4.2.1) applies and there is zero duality gap with the primal program. However, dual functionals feasible to (4.1.8) may no longer be countably additive. Let $\bar{\psi}$ be a dual functional over ℓ_∞ that satisfies $\bar{\psi}(\{a_i\}_{i=1}^\infty) := \lim_{i\to\infty} a_i$ for every convergent sequence $\{a_i\}_{i=1}^\infty$. Such a dual functional is guaranteed to exist (for details, see Lemmas 16.29 and 16.30 in Aliprantis and Border [2006]). Observe also that $\bar{\psi}$ is feasible to (4.1.8) and its dual objective value, $\bar{\psi}(\{-1,0,0,\dots\}) = 0$, is equal to the optimal value of the primal. Thus, $\bar{\psi}$ is an optimal dual functional and there is zero duality gap. $\triangleleft$

The optimality of $\bar{\psi}$ in the previous example highlights three serious issues that are not present in finite dimensions. First, $\bar{\psi}$ is difficult to characterize precisely. The only structure we have specified is that it is linear and acts like a "limit evaluator" on elements of ℓ_∞ that converge. Outside of the subspace of convergent sequences, little is known about how $\bar{\psi}$ operates.

Second, $\bar{\psi}$ fails countable additivity. Consider the sequence $(1,1,\dots)$ and observe that $1 = \bar{\psi}((1,1,\dots)) \neq \sum_{i=1}^\infty 1 \cdot \bar{\psi}(e_i) = 0$. This implies that we lose the familiar interpretation from finite dimensions that there exists a dual price $\bar{\psi}(e_i) = \bar{\psi}_i$ on the ith constraint so that if each constraint i is perturbed by sufficiently small ϵ_i, then the change in objective is $\sum_{i\in\mathbb{N}} \bar{\psi}_i \epsilon_i$.

Third, the dual (4.1.8) is not analogous to the finite linear programming dual. The constraints are not defined using the "transpose" of the original constraint matrix and instead are expressed in terms of dual functionals that are not necessarily countably additive.

Moreover, Example 4.1.2 illustrates how the choice of the vector space for the primal constraint space affects the structure and interpretation of the dual. The dual price vectors for (4.1.7) are easily characterized – they are finite support vectors. However, the interior of the positive cone of the constraint space is empty and there is a duality gap. Alternatively, choosing a constraint space with a positive cone that has a nonempty interior generates the dual program (4.1.8) with zero duality gap. However, the resulting structure of the dual problem and the structure of the optimal dual functionals are "undesirable."

In the finite case, dual functionals can always be interpreted as a price on each constraint; however Example 4.1.2 demonstrates this does not generalize to infinite dimensions. If countable additivity holds then this nice interpretation does carry over. In the general setting of Riesz spaces such "undesirable" dual functionals fail to satisfy the key property of *σ-order continuity*, which is equivalent to a generalized notion of countable additivity.

4.1.2 Our Contributions

Our main results show that the interplay between the existence of interior points and singular dual functionals observed in Example 4.1.2 is not an accident. Our results apply to constraint spaces that are infinite dimensional Riesz spaces. We make no topological assumptions and work with the algebraic notion of core points (defined in Section 4.2) rather than interior points. We show that if the positive cone of the constraint space has a core point then there is zero duality gap with the algebraic dual. Moreover, in a broad class of spaces – Riesz spaces that satisfy either σ-order completeness or the projection property – if the positive cone of the constraint space has a core point then singular dual functionals exist. We call this phenomenon the *Slater conundrum*. On the one hand, the existence of a core point ensures a zero duality gap (a desirable property), but on the other hand, existence of a core point implies the existence of singular dual functionals (an undesirable property).

This approach borrows concepts from geometric functional analysis (see for instance, Holmes [1975]) and connects them to concepts in Riesz space theory (see for instance, Aliprantis and Border [2006]). Proposition 4.2.5 unites the concept of a *core point* from geometric functional analysis to the concept of an *order unit* in Riesz spaces. This provides a bridge between two streams of literature that, to the authors' knowledge, produces a novel approach to the study of infinite dimensional optimization problems.

Corollary 4.2.2 provides a constraint qualification that relies on the existence of a core point in the positive cone of the constraint space. We call this an algebraic constraint qualification because the concept of core can be defined in any ordered vector space. The use of core points for constraint qualifications was first introduced by Rockafellar [1974]. However, the condition given applied to optimization problems over Banach spaces and their Fenchel duals. Our algebraic constraint qualification is an extension of these conditions to include general ordered vector spaces.

An advantage of the algebraic constraint qualification is that every interior point in a locally convex topological vector space is a core point (see Holmes [1975]). Therefore, if the algebraic constraint qualification fails to hold, then Slater's constraint qualification in any locally convex topology also fails. This property allows us to investigate which ordered vector spaces have interior points in their positive cones and determine the structure of dual functionals on these spaces.

A main result is Theorem 4.3.8, where we construct a singular dual functional from an order unit and a sequence of vectors that order converge to that order unit. This result is used to establish the existence of singular dual functionals in general Riesz spaces that need not possess any further topological or completeness properties required by other known results (including our own Theorem 4.3.13 and Theorem 4.5.13 (see also Wnuk [1999])). Next, we show (Theorem 4.3.13) that the general class of Riesz spaces with order units that are either σ-order complete or satisfy the projection property always have a singular dual functional in the algebraic dual space.

Riesz spaces are ordered vector spaces equipped with a lattice structure (see Section 4.2.2). We focus on Riesz spaces and not other classes of vector spaces for four reasons. First, an ordered vector space is necessary for constrained optimization. The common notion of being "constrained" is based on the concept of ordering. For instance, in finite-dimensional linear programming, inequality constraints $Ax \geq b$ and nonnegativity constraints $x \geq 0$ are defined by the standard ordering of finite-dimensional Euclidean space.

Second, the order and lattice assumptions endow an infinite dimensional vector space with just enough structure for familiar properties in $\mathbb{R}^n$ to have meaningful analogues. On a Riesz space we can define absolute value and order convergence. Also, the concept of disjointness (Section 4.2.2), generalizes the idea that a nonzero vector can have both zero and nonzero "components."

Third, the *order dual* of a Riesz space is the most general dual space that needs to be considered in optimization when the constraint space Y is an ordered vector space. Given vector space Y, the largest dual space is the algebraic dual Y'. However, when Y is a Riesz space, the order dual $Y^\sim$ consisting of dual functionals that are order-bounded is natural since it contains all of the positive linear functionals in Y'. By Proposition 4.2.4 the optimal value of the Lagrangian dual defined over Y' and the optimal value of the Lagrangian dual defined over $Y^\sim$ are equal. This implies that the order-dual structure of $Y^\sim$ can be used without loss of optimality when Y is a Riesz space.

Fourth, Riesz spaces provide the lattice structure which is needed to precisely classify which dual functionals are easy to characterize and interpret, and those that are not (i.e., those that are not countably additive) through the concept of order continuity. In an ordered vector space, the underlying algebraic structure is insufficient to separate countably additive dual functionals from those that are not countably additive.

An alternative to imposing a lattice structure on an ordered vector space is to endow it with a topology. This topological vector space approach is far more common in the optimization literature than the Riesz space approach. We show in Remark 4.3.4 that it is necessary to impose a lattice structure in order to distinguish the desirable from the undesirable dual functionals. However, because topological thinking is so pervasive, many researchers add the norm topology in addition to the lattice structure and work with Banach lattices (for example, Aliprantis and Burkinshaw [2006], Wnuk [1999], and Zaanen [1983]). A key contribution of this chapter is to show that the additional topological structure is not necessary for establishing fundamental results in optimization. Riesz spaces have the minimum structure necessary to characterize the Slater conundrum and establish the conundrum as a fundamental problem endemic to convex optimization. We adhere to the dictum expressed by Duffin and Karlovitz [1965] of "the desirability of omitting topological considerations" from a position of both enhanced clarity and enhanced generality. Indeed, one of our key results, Theorem 4.3.13, is more general than similar results obtained for Banach lattices. See Remark 4.3.15 and Theorem 4.5.13 in Appendix 4.5.4.

The main result of this chapter may be interpreted as being "negative" because we show that the Slater condition for zero duality gap implies the existence of singular functionals which are difficult to characterize. However, all is not lost when it comes to certain specially structured problems. In Section 4.4 we provide two sets of sufficient conditions that "resolve" the Slater condundrum for linear programs by guaranteeing the existence of optimal dual solutions with no singular component whenever the order dual is solvable. Previous studies, such as Ponstein [1981] and Shapiro [2005] proposed conditions in the special settings of ℓ_∞ and L_∞. Our conditions are stated over general Riesz spaces and generalize these approaches.

4.1.3 Literature Review

Our work is related to several streams of literature in economics and optimization. We briefly outline them here.

The Slater conundrum has been explored by others in specific contexts. Rockafellar [1974]; Rockafellar and Wets [1976b,a] and Ponstein [1981, 2004] observe that L_∞ is the only L_p space with an interior point in its positive cone. They further point out that algebraic dual of L_∞ contains dual functionals that are not countably additive. They conclude that the only L_p space where Slater's constraint qualification can be applied has dual functionals that are difficult to characterize and interpret. Our development avoids topological and measure-theoretic arguments (as used in the L_∞ case) and works in greater generality by focusing on primitive algebraic and order properties. This level of abstraction demonstrates that the Slater conundrum is endemic to infinite dimensional optimization at its very foundation.

In the stochastic programming literature, Rockafellar and Wets [1976a] emphasize the central role played by singular dual functionals in a complete duality theory for convex stochastic programs. Previously, authors ignored such dual functionals and worked only with countably additive dual functionals. This meant they had to accept the possibility of duality gaps. Although researchers begrudgingly accept that general optimality conditions involve singular dual functionals in an essential way, they are considered to be "unmanageable" from a practical perspective and fastidiously avoided; see for example, Rockafellar and Wets [1976b]. There has also been a substantial amount of research devoted to finding problem structures that do not have optimal singular dual functionals (see Dempster [1988] for a summary). In Ponstein [1981], Ponstein makes a careful study of singular dual functionals in general convex optimization problems with constraint spaces in L_∞. He gives conditions that justify ignoring singular dual functionals without loss of optimality, generalizing some well-known conditions in the stochastic programming literature. Ponstein emphasizes the computational intractability of singular dual functionals.

Other authors have also pursued an order-algebraic approach to optimization while eschewing topological concepts. Holmes's classic monograph Holmes [1975] on geometric functional analysis sets the stage for order-algebraic approaches by dedicating a large initial part of his monograph to an investigation of functional analysis without reference to topology. Similarly, Anderson and Nash [1987] ground their duality theory with a study of algebraic duality before introducing topological notions. In this context they establish weak duality and complementary slackness, as well as lay the algebraic foundation for an extension of the well-known simplex method to infinite dimensional linear programs. Following a similar approach, Shapiro [2005] considers conditions that imply zero duality gap for the algebraic dual of infinite conic programs and only later introduces topologies.

Our work, in some ways, parallels developments in theoretical economics on pricing and equilibria in infinite dimensional commodity spaces (see for instance, Aliprantis and Brown [1983]). These studies feature Riesz spaces arguments that are akin to ours. While the literature on Riesz spaces is quite extensive, with many quality texts ranging from the introductory level to the advanced (see Aliprantis and Border [2006], Aliprantis and Burkinshaw [2006], Luxemburg and Zaanen [1971] and Zaanen [1983, 1997]), Riesz spaces are rarely mentioned in optimization theory. A key to our approach is to develop novel results for Riesz spaces and apply these results to infinite dimensional optimization.

The Slater conundrum is driven by constraint qualifications for zero duality gap that require the existence of interior points in the positive cone in the constraint space. Such *interior point* constraint qualifications are not the only approaches to establishing zero duality gap results in infinite dimensional optimization. Researchers have long been aware of the limitations of interior point conditions in the L_p spaces. This awareness has motivated several alternate approaches that are worth briefly mentioning here.

First, some researchers have generalized the concept of an interior of a set. A powerful generalization, the quasi-relative interior, was introduced by Borwein and Lewis [1992].

Interior points are contained in the quasi-relative interior, as are core points Borwein and Lewis [1992]. For example, while the positive cone $\{x \in L_p[0,1] : x(t) \geq 0 \text{ a.a. } t \in [0,1]\}$ in $L_p[0,1]$ for $1 \leq p < \infty$ has no interior in its norm topology, the quasi-relative interior is $\{x \in L_p[0,1] : x(t) > 0 \text{ a.a. } t \in [0,1]\}$ Borwein and Lewis [1992]. Unfortunately, the direct extension of Slater's constraint qualification using the expanded set of quasi-relative interior points does not hold. The existence of a feasible point mapping to a quasi-relative interior point in a positive cone is not sufficient to establish zero duality gap without additional assumptions (see Appendix 4.5.2 for a concrete example). For instance, the constraint qualification presented in Borwein and Lewis [1992] requires that the constraint space be finite dimensional. A variety of constraint qualifications based on the quasi-relative interior were later introduced (see for example, Boţ [2010] and Grad [2010]). By considering the structure of the positive cone, rather than the topological structure of a vector space, our approach avoids the need for additional assumptions by focusing attention on vector spaces where core points exist.

Second, many researchers have focused on topological notions such as closeness, boundedness and compactness to drive duality results, rather than interior points Anderson and Nash [1987]; Boţ [2010]; Grinold [1968]. This approach is particularly useful when interior points are known not to exist. The drawback is that these conditions are generally thought to be difficult to verify in practice, despite their theoretical elegance.

Finally, others have established zero duality gap results using limiting arguments and finite approximations in combination with the basic duality results for finite dimensional linear programming. These include papers on separated continuous linear programs (Pullan [1993]) and countably infinite linear programs (Ghate and Smith [2013] and Romeijn et al. [1992]). These studies avoid certain topological arguments by leveraging duality results from finite dimensional linear programming and careful reasoning about limiting behavior. However, our approach is quite different. We consider general inequality-constrained convex programs (not just linear programs) and do not derive our results from duality results in the finite dimensional setting.

The remainder of the chapter is structured as follows. In Section 4.2.1 we give a constraint qualification that is sufficient for a zero duality gap between the primal and the algebraic dual in any ordered vector space. Section 4.2.2 is a very brief tutorial on Riesz spaces and provides the necessary background material. In Section 4.2.3, we show that the algebraic dual and order dual (a Riesz space concept) are equivalent on the positive cone and that a core point corresponds to an order unit (a Riesz space concept). The main results of the chapter are in Section 4.3 where we establish the Slater conundrum: that Slater points lead to bad (singular) dual functionals in the most general setting possible. Section 4.4 describes two sets of sufficient conditions for "working around" the Slater conundrum in specially structured problems.

We provide several appendices to supplement our results. Appendix 4.5.1 provides a proof of the results in Section 4.2.1. Appendix 4.5.2 discusses the differences between our algebraic constraint qualification and the constraint qualifications based on the quasi-relative interior. Appendix 4.5.3 includes definitions and results in Riesz space theory that are required in the proofs of some of our results. Appendix 4.5.4 examines the Slater conundrum from the viewpoint of Banach lattices.

4.2 Optimization in Ordered Vector Spaces

4.2.1 Lagrangian Duality

We first review some of the notation, definitions, and concepts of convex optimization and Lagrangian duality. Consider an *ordered vector space* Y containing a pointed, convex cone P. The cone P, called the *positive cone*, defines the vector space ordering $\succeq_P$, with $y \succeq_P \bar{y}$ iff $y - \bar{y} \in P$. The notation $y \succ_P \bar{y}$ indicates that $y - \bar{y} \in P$ and $y \neq \bar{y}$. Other authors have used $y \succ_P \theta_Y$ to mean that y lies in the interior of the cone P. We avoid this usage in favor of explictly stating when a vector lies in the interior of a set.

Consider the following inequality constrained convex program,

$$\begin{aligned} \inf \;\; & f(x) \\ \text{s.t.} \;\; & G(x) \preceq_P \theta_Y \\ & x \in \Omega, \end{aligned} \tag{CP}$$

where Ω is a convex set contained in a vector space X, $f : \Omega \to \mathbb{R}$ is a convex functional with domain Ω and $G : \Omega \to Y$ is a P-convex map from Ω into the ordered vector space Y and θ_Y represents the zero element of Y.

Let Y' denote *algebraic dual* of Y, set of linear functionals over Y. Take $\psi \in Y'$. The evaluation $\psi(y)$ of ψ at $y \in Y$ is alternatively denoted by $\langle y, \psi \rangle$; that is, $\langle y, \psi \rangle = \psi(y)$. We use the evaluation notation ($\psi(y)$) and the dual pairing notation ($\langle y, \psi \rangle$) interchangeably and favor the notation that lends the greatest clarity to a given expression. The *algebraic dual cone* of P is $P' := \{\psi \in Y' : \langle y, \psi \rangle \geq 0, \ \forall y \in P\}$ and the elements of P' are called *positive dual functionals* on Y. The restriction of the dual cone $P^{'}$ to a subspace W of Y', is denoted by $Q_W = P' \cap W$ and is called the *positive dual cone* with respect to W.

The Lagrangian function, $L : Y' \to \mathbb{R}$ for (CP) is $L(\psi) := \inf_{x \in \Omega} [f(x) + \langle G(x), \psi \rangle]$. Using this definition of $L(\psi)$, a family of dual programs for (CP) is derived as follows. Let

W be a subspace of the algebraic dual Y'. The Lagrangian dual program (D_W) of (CP) with respect to W is

$$\begin{aligned} &\sup \ L(\psi) \\ &\text{s.t. } \ \psi \in Q_W. \end{aligned} \tag{D_W}$$

The optimal value of an optimization problem $(\cdot)$ is denoted $v(\cdot)$. *Weak duality* holds when the value of the primal program is greater than or equal to the value of the dual. If $v(CP) = v(D_W)$, then the primal and dual programs have *zero duality gap*. As is well-known (see for instance Anderson and Nash [1987]) weak duality always holds for the Lagrangian dual program (D_W) regardless of the choice of W.

Slater's constraint qualification is perhaps the most well-known sufficient condition for zero duality gap between the primal program (CP) and its *topological dual program* (D_{Y^*}). When the constraint space Y is a locally convex topological vector space, its *topological dual* Y^* is the set of dual functionals that are *continuous* in the topology on Y. Slater's constraint qualification states that there is zero duality gap when $-G : \Omega \to Y$ maps a point in Ω to an interior point (in the topology that defines Y^*) of the positive cone P.

Theorem 4.2.1 (Slater's Constraint Qualification, Ponstein [2004], Theorem 3.11.2)**.** *Let Y be a locally convex topological vector space with positive cone P and topological dual Y^*. If there exists an $\bar{x} \in \Omega$ such that $-G(\bar{x}) \in \text{int}(P)$, then there is an optimal dual solution $\bar{\psi} \in Y^*$ and $v(CP) = v(D_{Y^*})$.*

The set of interior points, $\text{int}(P)$, obviously depends upon the selection of the locally convex topology. The proof of Theorem 4.2.1 uses the existence of an interior point to construct separating hyperplanes from dual funtionals that are continuous in the topology on Y. Therefore, one would like select a locally convex topology on Y such that: 1) *every* dual functional defined on Y is continuous so that the set of dual functionals is large for the goal of closing the duality gap (and thus the topological and algebraic dual are the same,

$Y' = Y^*$) and 2) the set $\text{int}(P)$ is the largest possible set of interior points. These two goals are achieved by using a locally convex topology defined by *core points*.

Given a vector space Y and a subset $A \subseteq Y$, a point $a \in A$ is a *core point* of A if for every $y \in Y$, there exists an $\epsilon > 0$ such that $a + \lambda y \in A$ for all $0 \leq \lambda \leq \epsilon$. The set of core points of A is denoted $\text{cor}(A)$.

Corollary 4.2.2 (Algebraic Constraint Qualification)**.** Let Y be a vector space with positive cone P and algebraic dual $Y^{'}$. If there exists an $\bar{x} \in \Omega$ such that $-G(\bar{x}) \in \text{cor}(P)$, then there is an optimal dual solution $\bar{\psi} \in Y^{'}$ and $v(CP) = v(D_{Y'})$.

Corollary 4.2.2 follows immediately from Theorem 4.2.1 and the fact that the core points can be used to define is a locally convex topology on Y where $\text{int}(A) = \text{cor}(A)$. For details, see Appendix 4.5.1. A core point is a purely geometric concept and the beauty of Corollary 4.2.2 is that it applies to any ordered vector space. This is in keeping with our philosophy to present results in the most general setting possible. Furthermore, all interior points in any locally convex topological vector space are core points; that is, for all subsets $A \subset Y$, $\text{int}(A) \subseteq \text{cor}(A)$ (see Holmes [1975]). If P does not have a core point, then P does not have an interior point in any locally convex topology. This implies that the existence of a core point in P is the most general Slater condition possible.

Remark 4.2.3.

Recently, authors have introduced constraint qualifications using quasi-relative interior points and other topological alternatives Borwein and Lewis [1992]; Boţ [2010]; Grad [2010]. Constraint qualifications based on a quasi-relative interior point require additional structure in order to prove that there is zero duality gap. In Appendix 4.5.2 we show that the existence of a quasi-relative interior point is not sufficient to guarantee a zero duality gap with the Lagrangian dual. $\triangleleft$

4.2.2 Riesz Spaces

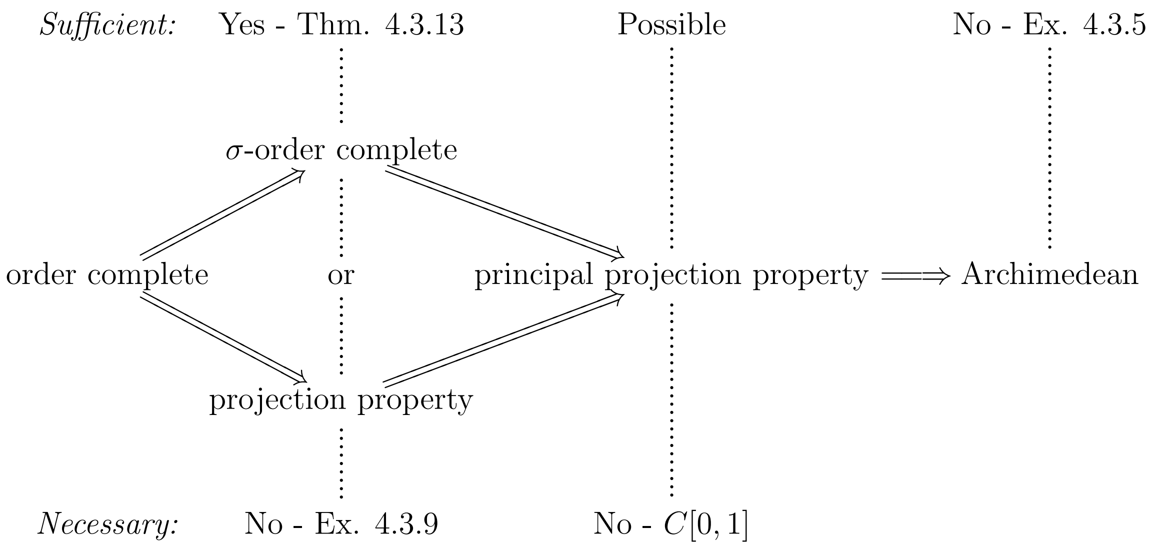

Figure 4.1: The Slater conundrum - necessary and sufficient conditions

We explore the effect that the existence of a core point in the positive cone has on the structure and interpretation of the dual functionals. For reasons outlined in the introduction, Riesz spaces provide a natural setting for this exploration.

We begin with a few basic definitions and concepts from Riesz space theory needed to understand the statements of our results. This includes definitions of the classifications of Riesz spaces in Figure 4.1. This figure gives the relationships among important classes of Riesz spaces. Luxemburg and Zaanen [1971] (Theorem 25.1) refer to these relationships as the *main inclusion theorem*. The main inclusion theorem is used to gain a precise understanding of our contributions in Section 4.3 (See Remark 4.3.14). Additional results on Riesz spaces used in proofs, but not needed to understand the statements of our results, are found in Appendix 4.5.3. For those interested in more details, Chapters 8 and 9 of Aliprantis and Border [2006] provide a thorough introduction to Riesz spaces.

An ordered vector space E is a Riesz space if the vector space is also a lattice; that is, each pair of vectors $x, y \in E$ has a supremum denoted $x \vee y$ and an infimum denoted $x \wedge y$.

Common examples of Riesz spaces are in Example 8.1 in Aliprantis and Border [2006]) and include the L_p spaces and spaces of continuous functions. Let E be a Riesz space ordered by the positive cone E^+. Let θ_E denote the zero vector of E. For each vector $x \in E$, the positive part x^+, the negative part x^- and the absolute value $|x|$ are $x^+ := x \vee \theta_E$, $x^- := x \wedge \theta_E$ and $|x| := x \vee (-x)$. Two vectors $x, y \in E$ are *disjoint*, denoted $x \perp y$, whenever $|x| \wedge |y| = \theta_E$. A set of vectors are *pairwise disjoint* if each pair of distinct vectors are disjoint. Given a subset of a Riesz space $S \subseteq E$, its *disjoint complement* is $S^d := \{y \in E : x \perp y \text{ for all } x \in S\}$. A set S is order bounded from above if there exists an upper bound $u \in E$, such that $x \preceq u$ for all $x \in S$. Similarly, a set is order bounded from below if there exists a lower bound and is *order bounded* if the set is order bounded both from above and below. A set S is *solid* when $|y| \preceq |x|$ and $x \in S$ imply that $y \in S$.

A Riesz space E is *order complete* (sometimes called Dedekind complete) if every non-empty subset of E that is order bounded from above has a supremum. Similarly, a Riesz space is *σ-order complete* if every countable subset with an upper bound has a supremum. The L_p spaces for $1 \leq p < \infty$ are σ-order complete.

Let $\{x_\alpha\}$ be a net of vectors in E and let $\{x_n\}$ represent a sequence. The notation $x_\alpha \uparrow\preceq x$ means that $\{x_\alpha\}$ is order bounded from above by x. When $\sup x_\alpha = x$, we write $x_\alpha \uparrow x$. Define $x_\alpha \downarrow x$ similarly. A Riesz space E is *Archimedean* if $\frac{1}{n}x \downarrow \theta$ for each $x \succeq \theta_E$.

A net $\{x_\alpha\}$ in a Riesz space E *converges in order* to x denoted as $x_\alpha \xrightarrow{\circ} x$ if and only if there exists a net $\{y_\alpha\}$ with the same directed set such that $|x_\alpha - x| \preceq y_\alpha$ for each α and $y_\alpha \downarrow \theta_E$. Note that the notion of order convergence involves the absolute value, and thus cannot be defined in an ordered vector space without a lattice structure. A subset S in E is *order closed* if for any net in S with $x_\alpha \xrightarrow{\circ} x \in E$ has $x \in S$.

A vector subspace of a Riesz space E is a *Riesz subspace* if it is closed under the lattice operations of E. A solid Riesz subspace is called an *ideal.* A *principal ideal* $E_x \subseteq E$ is an ideal generated by a vector $x \in E$ and is defined as $E_x := \{y \in E : \exists \lambda > 0 \text{ s.t. } |y| \preceq \lambda |x|\}$.

A band is an ideal that is order closed. The band B_x which consists of the order closure of E_x is called the *principal band* generated by $x \in E$. A band B is called a *projection band* if $E = B \oplus B^d$ where $\oplus$ denotes a direct sum of vector subspaces, meaning that every element x of E can be written uniquely as $x = y + z$ with $y \in B$ and $z \in B^d$ and $|y| \wedge |z| = \theta$. A Riesz space E has the *projection property* if every band is a projection band. A Riesz space E has the principal projection property if every principal band is a projection band. Let B be a projection band in a Riesz Space E. By definition of projection band, $E = B \oplus B^d$ and for every $x \in E$ there exists an $x_1 \in B$ and an $x_2 \in B^d$ such that $x = x_1 + x_2$. Let $P_B : E \to B$ be defined as $P_B(x) := x_1$.

A dual functional $\psi : E \to \mathbb{R}$ is *order bounded* if it maps order bounded sets in E to order bounded sets in $\mathbb{R}$. The set of all order bounded dual functionals on a Riesz space E is called the *order dual* of E and is denoted $E^{\sim}$. A dual functional $\psi \in E^{\sim}$ on a Riesz space E is *order continuous* if $\psi(x_\alpha) \to 0$ for all nets $\{x_\alpha\}$ that order converge to θ_E. A dual functional $\psi \in E^{\sim}$ on a Riesz space E is *σ-order continuous* if $\psi(x_n) \to 0$ for all sequences $\{x_n\}$ that order converge to θ_E. The set of dual functionals that are σ-order continuous $E_c^{\sim}$ form a subspace of the order dual called the *σ-order continuous dual.* The order dual can be expressed as the direct sum of $E_c^{\sim}$ and its complementary disjoint subspace $E_s^{\sim} := (E_c^{\sim})^d$; that is, $E_c^{\sim} \oplus E_s^{\sim} = E^{\sim}$ (see Theorem 8.28 of Aliprantis and Border [2006]). The dual functionals $\psi \in E_s^{\sim}$ are called *singular dual functionals.*

4.2.3 An Order-Algebraic Approach

Consider the convex, inequality constrained program (CP) where the constraint space Y is an ordered vector space. When Y is also a Riesz space, then its order dual $Y^{\sim}$ generates the Lagrangian dual program $(D_{Y^{\sim}})$. The added structure of the order dual allows us to further characterize and interpret dual functionals. Corollary 4.2.2 applies only to the algebraic dual program. In this section, we show this result also applies to the order dual program. The resulting order-algebraic approach connects the theory developed in general ordered vector spaces to the structure and interpretations available in Riesz spaces.

By definition, the order dual of a Riesz space Y contains only order bounded dual functionals and is a subset of the algebraic dual. The feasible region of the algebraic dual program $(D_{Y'})$ consists of positive dual functionals. The following proposition shows that the order dual program and algebraic dual program have the same feasible region and are therefore equivalent.

Proposition 4.2.4. *For the convex program* (CP)*, the value of the algebraic dual equals the value of the order dual; that is,* $v(D_{Y'}) = v(D_{Y^\sim})$.

Proof. It suffices to show that the feasible regions of $(D_{Y'})$ and $(D_{Y^\sim})$ are equal; that is, $P' = Q_{Y^\sim}$ where $Q_{Y^\sim} = P' \cap Y^\sim$. Clearly $Q_{Y^\sim} \subseteq P'$ so it suffices to show $Q_{Y^\sim} \supseteq P'$. Let $\psi \in P'$, that is ψ is a positive dual functional over Y. Let $x, y, z \in Y$ and $x \preceq y \preceq z$. Then, $\psi(y - x) \geq 0$ and $\psi(z - y) \geq 0$. This implies $\psi(x) \leq \psi(y) \leq \psi(z)$ and the dual functional ψ is order bounded. That is, $\psi \in Y^\sim$. This implies $\psi \in P' \cap Y^\sim = Q_{Y^\sim}$. We conclude $P' = Q_{Y^\sim}$. □

Proposition 4.2.4 provides a crucial link between the algebraic constraint qualification and the order dual. Restricting the space of all dual functionals to the order dual does not create a duality gap with the primal. Therefore, if a core point satisfies the algebraic constraint qualification then $v(CP) = v(D_{Y'})$ and Proposition 4.2.4 implies zero duality gap between the primal and the order dual.

Core points of the positive cone relate to a fundamental concept in the theory of Riesz spaces called order units. An element $e \succ \theta_E$ in a Riesz space E is an *order unit* if for each $x \in E$ there exists a $\lambda > 0$ such that $|x| \leq \lambda e$. The equivalence between order units and core points is given below. This result was also known by Aliprantis and Tourky [2007]. However, Aliprantis and Tourky do not use this result within the context of optimality conditions. We provide a proof for the sake of completeness.

Proposition 4.2.5. *If E is a Riesz space and $e \succ \theta_E$, then e is an order unit of E if and only if e is a core point of the positive cone E^+.*

Proof. Assume that e is an order unit and let z be an an arbitrary element of E. We show there exists an $\epsilon > 0$ such that $e + \lambda z \in E^+$ for all $\lambda \in [0, \epsilon]$. If e is an order unit then by definition there exists an $\alpha > 0$, such that $|z| \preceq \alpha e$. Let $\epsilon = 1/\alpha$. Then $\lambda|z| \preceq e$ for all $\lambda \in [0, \epsilon]$. This implies $-\lambda z \preceq e$ and therefore $\theta_E \preceq e + \lambda z$. Therefore, $e + \lambda z \in E^+$ for all $\lambda \in [0, \epsilon]$ and by definition, e is a core point of E^+. Next, assume e is a core point of E^+ and let z be any element in E. Then there exists a $\epsilon_+ > 0$ such that $e + \lambda z \in E^+$ for all $\lambda_+ \in [0, \epsilon_+]$. If $e + \lambda_+ z \in E^+$, then $e + \lambda_+ z \succeq \theta_E$ and therefore $-z \preceq \frac{1}{\lambda_+} e$. Applying the same logic to $-z$ there exists a λ_- that gives $z \preceq \frac{1}{\lambda_-} e$. Take $\lambda = \min\{\lambda_+, \lambda_-\}$. Then $-z \preceq \frac{1}{\lambda} e$ and $z \preceq \frac{1}{\lambda} e$ implies $-z \vee z \preceq \frac{1}{\lambda} e$. By definition $|z| = z \vee -z$ so $|z| \preceq \frac{1}{\lambda} e$ and e is an order unit of E. □

Proposition 4.2.4 and Proposition 4.2.5 provide the foundations for our order-algebraic approach. Since the value of the order dual and the algebraic dual are equal, the algebraic constraint qualification provides a sufficient condition for zero duality gap between the primal program and the order dual. In Riesz spaces, core points of the positive cone are order units. Therefore, only Riesz spaces that contain order units can satisfy Corollary 4.2.2.

4.3 The Slater Conundrum

This section contains the main results of the chapter. We show for a broad class of Riesz spaces (σ-order complete spaces) that if the positive cone has a core point, then it is necessary to either show that the optimal dual functional does not have a singular component, or somehow characterize the singular components. To our knowledge, this has never been shown in the broad contexts of Theorems 4.3.8, 4.3.12, and 4.3.13.

4.3.1 Duality and Countable Additivity

In this subsection we expand a theme in Example 4.1.2 regarding the undesirability of dual functionals that are not countably additive. Countable additivity (defined in Example 4.1.2 for sequence spaces) extends the familiar interpretations of dual functionals to infinite dimensional Riesz spaces. In fact, the σ-order continuous dual functionals introduced in Section 4.2.2 are countably additive. We now make precise the concept of countable additivity for arbitrary Riesz spaces.

Given any sequence $\{a_i\}_{i=1}^{n}$ of real numbers, the limit of partial sums $\sum_{i=1}^{n} a_i$ is written as $\sum_{i=1}^{\infty} a_i$ whenever the limit of partial sums exists. Loosely speaking, the reason a lattice structure is added to an ordered vector space is to give the corresponding Riesz space properties that mimic the real numbers as closely as possible. Given a lattice structure, order convergence can be defined, and it is similar to convergence of real numbers. Now assume that the a_i are vectors in an arbitrary Riesz space E. If $x_n \stackrel{\circ}{\to} \bar{x}$ where x_n is the partial sum $x_n = \sum_{i=1}^{n} a_i$ we follow the common practice used for real numbers and write $\bar{x} := \sum_{i=1}^{\infty} a_i$.

Let ψ be a σ-order continuous dual functional defined on E. If $x_n \stackrel{\circ}{\to} \bar{x}$, then $(x_n - \bar{x}) \stackrel{\circ}{\to} \theta_E$ and it follows from the definition of σ-order continuity that

$$\psi\left(\sum_{i=1}^{\infty} a_i\right) = \psi(\bar{x}) = \lim_{n\to\infty} \psi(x_n) = \lim_{n\to\infty} \psi\left(\sum_{i=1}^{n} a_i\right) = \lim_{n\to\infty} \sum_{i=1}^{n} \psi(a_i) = \sum_{i=1}^{\infty} \psi(a_i)$$

where the last equality follows from the fact that the sequence of partial sums of real numbers, $\sum_{i=1}^{n} \psi(a_i)$, converges to a real number $\psi(\bar{x})$. Since $\psi\left(\sum_{i=1}^{\infty} a_i\right) = \sum_{i=1}^{\infty} \psi(a_i)$ we say that ψ is *countably additive.*

Likewise, if a dual functional is countably additive it is σ-order continuous. Assume ψ is a countably additive dual functional on a Riesz space E and $\{x_n\}_{n=1}^{\infty}$ is a sequence in E such that $x_n \stackrel{\circ}{\to} \theta_E$. Define a new sequence $\{a_i\}_{i=1}^{\infty}$ from $\{x_i\}_{i=1}^{\infty}$ by $a_1 = x_1$ and $a_i = x_i - x_{i-1}$ for $i \geq 2$. Then $x_n = \sum_{i=1}^{n} a_i$ and $x_n \stackrel{\circ}{\to} \theta_E$ implies that the sequence of partial sums $\sum_{i=1}^{n} a_i$

order converges to θ_E. Since ψ is countably additive $\psi\left(\sum_{i=1}^{\infty} a_i\right) = \sum_{i=1}^{\infty} \psi(a_i) = \psi(\theta_E)$ and this implies

$$\lim_{n\to\infty} \psi(x_n) = \lim_{n\to\infty} \psi(\sum_{i=1}^{n} a_i) = \lim_{n\to\infty} \sum_{i=1}^{n} \psi(a_i) = \sum_{i=1}^{\infty} \psi(a_i) = \psi(\sum_{i=1}^{\infty} a_i) = \psi(\theta_E).$$

Then $\lim_{n\to\infty} \psi(x_n) = \psi(\theta_E)$ and ψ is σ-order continuous. Therefore countable additivity and σ-order continuity are equivalent and we have shown:

Proposition 4.3.1. *If ψ is a dual functional on E, then ψ is σ-order continuous if and only if it is countably additive.*

The next result shows all the dual functionals in the order dual $E^\sim$ are countably additive if and only if the order dual does not contain any nonzero singular dual functionals. Recall from Section 4.2.2 that a dual functional ψ in the order dual $E^\sim$ of Riesz space E can be written as the sum of a σ-order continuous dual functional $\psi_c \in E_c^\sim$ and a singular dual functional $\psi_s \in E_s^\sim$; that is, $\psi = \psi_c + \psi_s$ where ψ_c and ψ_s are unique.

Theorem 4.3.2. *Let E be a Riesz space. Then, $E^\sim$ contains a dual functional that is not countably additive if and only if $E^\sim$ contains a nonzero singular dual functional.*

Proof. ($\Leftarrow$) Suppose $E^\sim$ contains a nonzero singular dual functional $\psi \in E_s^\sim$. Since $E_s^\sim$ and $E_c^\sim$ are orthogonal, $E_c^\sim \cap E_s^\sim = \{\theta_{E'}\}$ and so $\psi \notin E_c^\sim$. Then by Proposition 4.3.1 ψ is not countably additive. ($\Rightarrow$) Suppose $E^\sim$ contains a dual functional ψ that is not countably additive. By Theorem 8.28 in Aliprantis and Border [2006], $\psi = \psi_c + \psi_s$ for $\psi_c \in E_c^\sim$ and $\psi_s \in E_s^\sim$. Then $\psi \notin E_c^\sim$ implies $\psi_s \neq \theta_{E'}$ and $E_s^\sim$ contains a nonzero element. □

By Theorem 4.3.2 singular dual functionals are clearly a problem because their existence implies that there are dual functionals in the order dual that are not countably additive. We show in Section 4.3.2 a tight connection between Riesz spaces E with order units and the order duals $E^\sim$ that have singular dual functionals.

Countable additivity allows us, in many cases, to write an infinite dimensional dual that is analogous to the finite dimensional dual. Consider the special case of *linear programs*. In linear programs with infinite dimensional constraint spaces, σ-order continuity plays an important role in *expressing* the dual in a familiar way. Consider the linear program

$$\begin{aligned} \inf_{x \in X} \quad & \varphi(x) \\ \text{s.t.} \quad & A(x) \succeq_P b \end{aligned} \tag{4.3.1}$$

where X is a vector space, $b \in Y$ where Y is an ordered vector space with positive cone P, $A : X \to Y$ is a linear map, and φ is a linear functional on X. The algebraic dual (see for instance Anderson and Nash [1987] for details) of (4.3.1) is

$$\begin{aligned} \sup_{\psi \in Y'} \quad & \psi(b) \\ \text{s.t.} \quad & A'(\psi) = \varphi \\ & \psi \in P', \end{aligned} \tag{4.3.2}$$

where $A' : Y' \to X'$ is the *algebraic adjoint* of A, defined by $\langle x, A'(\psi)\rangle = \langle A(x), \psi\rangle$. Giving a concrete expression for the adjoint A' is, in general, difficult. However, when the vector spaces are finite dimensional characterizing A' is easy.

If $X = \mathbb{R}^n$ and $Y = \mathbb{R}^m$ then (4.3.1) and (4.3.2) are easily expressed in terms of matrices. The linear map A is characterized by an m by n matrix (abusing notation in the standard way) $A = (a_{ij})$ where $a_{ij} \in \mathbb{R}$ for $i = 1, \ldots, m$ and $j = 1, \ldots, n$ and given $x \in \mathbb{R}^n$, $A(x) = \sum_{j=1}^{n} a_{\cdot j} x_j$ where $a_{\cdot j}$ is the jth column of matrix A. The algebraic adjoint map is characterized by the matrix transpose $A^\top$. To see this, recall that ψ is countably additive and characterized by the vector $\psi = (\psi_i)_{i=1}^{m}$ with $\psi_i \in \mathbb{R}$, where for $y \in \mathbb{R}^m$, $\langle y, \psi\rangle = \sum_{i=1}^{m} y_i \psi_i$.

Then

$$\begin{aligned}\langle A(x),\psi\rangle &= \sum_{i=1}^{m} A(x)_i\psi_i \\ &= \sum_{i=1}^{m}\left(\sum_{j=1}^{n} a_{ij}x_j\right)\psi_i = \sum_{j=1}^{n}\left(\sum_{i=1}^{m} a_{ij}\psi_i\right)x_j = \langle x, A'(\psi)\rangle. \qquad (4.3.3)\end{aligned}$$

Thus the adjoint operator corresponds to the usual matrix transpose, i.e. $\langle x, A'(\psi)\rangle = A^\top\psi(x)$.

The above analysis depends on two fundamental properties: (i) that the dual functional ψ over $\mathbb{R}^m$ is expressed as a real vector ψ in $\mathbb{R}^m$; and (ii) that it is permissible to swap the finite sums in (4.3.3).

Now consider $X = \mathbb{R}^n$ and $Y \subseteq \mathbb{R}^{\mathbb{N}}$ as in Example 4.1.2 where countable additivity is

$$\psi\left(\{a_i\}_{i=1}^{\infty}\right) = \psi\left(\sum_{i=1}^{\infty} a_i e_i\right) = \sum_{i=1}^{\infty}\psi\left(a_i e_i\right) = \sum_{i=1}^{\infty} a_i\psi\left(e_i\right) = \sum_{i=1}^{\infty} a_i\psi_i. \qquad (4.3.4)$$

Combining countable additivity of ψ with the fact X is an n-dimensional vector space gives

$$\begin{aligned}\langle A(x),\psi\rangle &= \sum_{i=1}^{\infty} A(x)_i\psi_i \\ &= \sum_{i=1}^{\infty}\left(\sum_{j=1}^{n} a_{ij}x_j\right)\psi_i = \sum_{j=1}^{n}\left(\sum_{i=1}^{\infty} a_{ij}\psi_i\right)x_j = \langle x, A'(\psi)\rangle \qquad (4.3.5)\end{aligned}$$

and once again, even though A has an infinite number of rows, we can write the adjoint as $A' = A^\top$. This phenomenon was demonstrated in Example 4.1.2 where the dual program (4.1.7) is characterized in the familiar way based on the "transpose" of the primal constraints. It is the absence of singular functionals in $Y^\sim$ that allows one to write a dual with this convenient dual representation. However, in (4.1.8) the dual program did not have a nice characterization because $\bar{\psi}$ was not countably additive.

When both X and Y are infinite dimensional sequence spaces countable additivity is not sufficient to justify writing the adjoint as the transpose matrix. Additional properties on the linear map A are required to apply Fubini's theorem and change the order of summation, as was done in (4.3.3). Identifying such additional conditions in the case of countably infinite linear programs, where X and Y are both sequence spaces, has been the focus of several studies including Ponstein [1981], Romeijn et al. [1992] and Ghate [2006].

Even in many function spaces countable additivity provides a structure for dual functionals that is analogous to the finite dimensional case. Consider an arbitrary $\sigma-$finite measure space (T, Σ, μ). The set of all real $\mu-$measurable functions is the Riesz space $M(T, \Sigma, \mu)$. As with the L_p spaces, we define two functions $f, g \in M(T, \Sigma, \mu)$ to be equivalent if they are equal everywhere except a set of measure zero. An *ideal* of a Riesz space E is a solid Riesz subspace of E. The following theorem shows that for any ideal of $M(T, \Sigma, \mu)$, a dual functional that is σ-order continuous has an appealing structure.

Theorem 4.3.3 (Zaanen [1983], Theorem 86.3)**.** *For any σ-order continuous dual functional ψ on an ideal F of $M(T, \Sigma, \mu)$, there exists a μ-measurable function p on T such that $\psi(y) = \int_{t \in T} p(t)y(t)d\mu$ holds for all $y \in F$. Furthermore, p is $\mu-$almost everywhere uniquely determined.*

Among the ideals of $M(T, \Sigma, \mu)$ are the associated $L_p(T, \Sigma, \mu)$ spaces for $1 \leq p \leq \infty$. Theorem 4.3.3 provides a large class of Riesz spaces where dual functionals that are σ-order continuous can be characterized by a measurable function p. This function is analogous to the vector $(\psi_i : i \in \mathbb{N})$ in Example 4.1.2. It assigns a value $p(t)$ to every element $t \in T$, which is interpreted analogously to a shadow price or marginal value. Furthermore, the integral structure in Theorem 4.3.3 is a convenient representation of the dual functional ψ and aids in expressing the adjoint provided additional properties allow for Fubini's theorem to apply.

Remark 4.3.4.

One might consider adding a topology to an ordered vector space rather than a lattice. However, by Proposition 4.3.1 countable additivity and order continuity are equivalent and a lattice structure is necessary to define order continuity and therefore distinguish the countable dual functionals from those that are not. The continuous dual functionals defined by a topology *are insufficient to distinguish countable additivity.* For example, by Theorem 4.5.9, when E is a Banach lattice – a Riesz space with a complete norm that is compatible with the lattice structure – the set of norm-continuous dual functionals is the order dual $E^{\sim}$. However, not every element of $E^{\sim}$ is σ-order continuous as we have seen, for example, in the case of ℓ_{∞}. $\triangleleft$

4.3.2 Order units and the existence of singular dual functionals

In the analysis that follows, we connect the theory of order units and singular dual functionals in Riesz spaces. Recall that order units (via Proposition 4.2.5) are fundamental to optimization due to their connection to positive core points and the algebraic constraint qualification (Corollary 4.2.2).

A simple conjecture is that in an infinite dimensional Riesz space, the existence of an order unit implies the existence of a singular dual functional in the order dual. The following examples demonstrates that this is not the case.

Example 4.3.5 (Luxemburg and Zaanen [1971], Example (v) on page 141 and Zaanen [1983], Example 103.5)**.**

Consider the Archimedean Riesz space E of all real functions f on an uncountable set T for which there exists a finite number $f(\infty)$ such that, given any $\epsilon > 0$, we have $|f(t) - f(\infty)| > \epsilon$ for at most finitely many $t \in T$. The function $e(t) = 1$ for all $t \in T$ is an order unit of E. However, Zaanen [1983] shows that all dual functionals in the order dual of E are σ-order continuous. $\triangleleft$

This example demonstrates the need for additional structure to guarantee the existence of a singular functional. The next example provides insight into structure leveraged in later proofs.

Example 4.3.6.

The space of bounded sequences ℓ_∞ has the order unit $e = (1, 1, 1, \ldots)$. Let $e_i \in \ell_\infty$ have 1 in component i and zero otherwise. Furthermore, define $x_n := \sum_{i=1}^{n} e_i$. Clearly, $x_n \xrightarrow{\circ} e$. Consider the positive dual functional $\bar{\psi}$ defined in Example 4.1.2. Then $\bar{\psi}(x_n) = 0$ for all n and $\bar{\psi}(e) = 1$. Thus $\bar{\psi}(e)$ is not σ-order continuous and $\bar{\psi} = \bar{\psi}_c + \bar{\psi}_s$ where $\bar{\psi}_s$ is a nonzero singular dual functional in ℓ_∞. $\triangleleft$

This example guides our approach to constructing singular dual functionals from order units in Theorem 4.3.8 below. The idea is to find a sequence of vectors $\{x_n\}_{n=1}^{\infty}$ where none of the x_n are order units, but $\{x_n\}_{n=1}^{\infty}$ order converges to an order unit $e \succ \theta_E$. The following result on extending dual functionals is used in the argument.

Theorem 4.3.7 (Krein-Rutman Theorem, Holmes [1975], Theorem 6B)**.** *Let X be an ordered vector space with positive cone P and let M be a subspace of X. Assume that $P \cap M$ contains a core point of P. Then any positive linear dual functional ψ on M admits a positive extension to all of X.*

Theorem 4.3.8. *Let E be an infinite dimensional Riesz space with order unit $e \succ \theta_E$. If there exists an increasing sequence $\{x_n\}$ of non-order units such that $x_n \xrightarrow{\circ} e$, then there exists a positive, nonzero, singular dual functional on E.*

Proof. Without loss of generality, assume that $x_n \succeq \theta_E$ for all n. Otherwise, replace x_n with the sequence $x_n^+ = x_n \vee \theta_E$. None of the x_n^+ are order units and by Lemma 8.15(ii) in Aliprantis and Border [2006], $x_n^+ \xrightarrow{\circ} e$. Let M be a subspace of E defined by $M := \text{span}(\{e\} \cup \{x_n\})$. Then every $y \in M$ is represented by a finite set of scalars $\{\lambda_n\}_{n=0}^{N}$ such that $y = \lambda_0 e + \sum_{n=1}^{N} \lambda_n x_n$. For shorthand, let y_x denote $\sum_{n=1}^{N} \lambda_n x_n$.

Claim 3. *For every* $y = \lambda_0 e + \sum_{n=1}^{N} \lambda_n x_n$ *in* M *the value of* λ_0 *in its representation is uniquely determined.*

Proof of Claim 3: We show that the order unit e cannot be represented as a linear combination of the $\{x_n\}$. Assume the opposite; that is, there exists a finite set of scalers $\{\mu_n\}_{n=1}^{N}$ such that $\sum_{n=1}^{N} \mu_n x_n = e$. Let $\overline{\mu} = \max\{|\mu_1|, |\mu_2|, \ldots, |\mu_N|\}$. Since $\{x_n\}$ is an increasing, nonnegative sequence, $\overline{\mu} N x_N \succeq \sum_{n=1}^{N} \mu_n x_n = e$. Next, consider an arbitrary $z \in E$. Since e is an order unit, there exists an $\alpha > 0$ such that $|z| \preceq \alpha e \preceq \alpha \overline{\mu} N x_N$. By definition, this implies that x_N is an order unit, arriving at a contradiction. Therefore, the value of λ_0 in the decomposition of y is uniquely determined. †

Define the functional $\psi_M : M \to \mathbb{R}$ by $\psi_M(y) = \psi_M(\lambda_0 e + y_x) := \lambda_0$ for all $y \in M$. This functional is well-defined by Claim 3. The following claim establishes the properties on ψ_M needed to apply the Krein-Rutman Theorem.

Claim 4. *The functional* ψ_M *is a positive linear functional on* M.

Proof of Claim 4: Let $y = \lambda_{0y} e + y_x$ and $z = \lambda_{0z} e + z_x$ be vectors in M and let $\alpha, \beta \in \mathbb{R}$. Then $\psi_M(\alpha y + \beta z) = \psi_M((\alpha\lambda_{0y} + \beta\lambda_{0z})e + \alpha y_x + \beta z_x) = \alpha\lambda_{0y} + \beta\lambda_{0z} = \alpha\psi_M(y) + \beta\psi_M(z)$. Hence, ψ_M is a linear functional on M. Next show that ψ_M is a positive functional. Assume otherwise. Then there exists a $y \in E$ such that $y = \lambda_0 e + \sum_{n=1}^{N} \lambda_n x_n$ where $y \succeq \theta_E$ and $\lambda_0 < 0$. Without loss, scale y such that $\lambda_0 = -1$. Then, $\sum_{n=1}^{N} \lambda_n x_n \succeq e$. As in the proof of Claim 3, let $\overline{\lambda} = \max\{|\lambda_1|, |\lambda_2|, \ldots, |\lambda_N|\}$. Then, since $\{x_n\}$ is an increasing, nonnegative sequence, $\overline{\lambda} N x_N \succeq e$. However, this implies that x_N is an order unit and yields a contradiction. †

Since $e \in E^+ \cap M$ is an order unit of E, it is a core point of E^+ by Proposition 4.2.5. Furthermore, Claim 4 states that ψ_M is a positive linear functional on M. Therefore, by the Krein-Rutman Theorem, ψ_M on M extends to a positive dual functional ψ on E. Notice that ψ is not σ-order continuous, since $x_n \overset{\circ}{\to} e$, $\psi(x_n) \to 0$ since $\psi(x_n) = \psi_M(x_n) = 0$, and $\psi(e) = 1$. Therefore, $\psi = \psi_c + \psi_s$ with $\psi_s \neq \theta_E$, since $E^\sim = E_c^\sim \oplus E_s^\sim$. Since $E^\sim$ is a Riesz space (see Theorem 8.24 in Aliprantis and Border [2006]) and ψ is positive, Theorem 4.5.6 implies that ψ_s is a positive singular dual functional on E. $\square$

Theorem 4.3.8 applies to arbitrary infinite dimensional Riesz spaces. To leverage this result, additional structure is added to a Riesz space to guarantee the existence of an order convergent sequence that satisfies the premise of Theorem 4.3.8. Below, in Theorem 4.3.13, we prove that all infinite dimensional Riesz spaces with order units that are either σ-order complete or have the projection property Riesz have sufficient additional structure to meet this criterion. However, as the following example demonstrates, this additional structure is not necessary for the existence of singular dual functionals.

Example 4.3.9 (Luxemburg and Zaanen [1971], Example (iv) page 140)**.**

The Riesz space $\ell^0_\infty(\mathbb{N})$ is the space of all sequences that are constant except for a finite number of components. For example, the sequence $(2, 0, 3, 1, 1, 1 \ldots)$ is an element of $\ell^0_\infty(\mathbb{N})$. Let $e = (1, 1, \ldots)$ be the vector of all ones. Then the set $\{e_i : i = 1, 2, 3, \ldots\} \cup \{e\}$ forms a Hamel basis of $\ell^0_\infty(\mathbb{N})$ and every dual functional can be characterized by assigning values to each vector in the set. The dual functionals on $\ell^0_\infty(\mathbb{N})$ that are not σ-order continuous are those that assign a value of zero to all e_i and a nonzero value to e. Let ψ be such a dual functional. Then $y_n \xrightarrow{o} e$ where $y_n = \sum_{i=1}^{n} e_i$, but $\lim_{n\to\infty} \psi(y_n) = 0 \neq \psi(e)$ and so $\ell^0_\infty(\mathbb{N})$ contains a singular dual functional. However, $\ell^0_\infty(\mathbb{N})$ is not σ-order complete and does not have the projection property. First consider $x_1 = (1, 0, 0, \ldots)$, $x_2 = (1, 1/2, 0, 0, \ldots)$, $x_3 = (1, 1/2, 1/3, 0, 0, \ldots)$, ... The sequence $\{x_n\}$ is order bounded from above by $e \in \ell^0_\infty(\mathbb{N})$ but clearly has no supremum in $\ell^0_\infty(\mathbb{N})$ and is therefore not $\sigma-$order complete. Next consider the band $B = \{x \in \ell^0_\infty(\mathbb{N}) : x(i) = 0, i\,\text{odd}\}$. Then $x \in B$ implies x has a finite number of nonzero even components. Clearly all of the $y \in B^d$ have the property that the even components of y must be zero and this implies every $y \in B^d$ has a finite number of nonzero odd components. Then $e \notin B \oplus B^d$ so B cannot be a projection band. $\triangleleft$

The proof of Theorem 4.3.12 shows how to construct, in any infinite dimensional Riesz space with order unit that is either σ-order complete or has the projection property, an increasing sequence $\{x_n\}$ of non-order units such that $x_n \xrightarrow{\circ} e$. Constructing $\{x_n\}$ requires the following lemma, due to Luxemburg and Zaanen [1971], used to generate an initial sequence from which we construct $\{x_n\}$.

Lemma 4.3.10 (Luxemburg and Zaanen [1971], Proposition 26.10)**.** Every infinite dimensional Archimedean Riesz space E contains an infinite set of pairwise disjoint elements.

The next example uses ℓ_∞ to motivate the steps of the proof of our key result Theorem 4.3.12.

Example 4.3.11.

Let E be the space of bounded sequences ℓ_∞ with the order unit $e = (1, 1, 1, \ldots)$. Assume Lemma 4.3.10 generates a sequence of vectors $\{u_n\}$ where for vector u_n, component $2n - 1$ is $(1/2)^n$ and the other components are zero. The first three vectors in $\{u_n\}$ are $u_1 = (\frac{1}{2}, 0, \ldots)$, $u_2 = (0, 0, \frac{1}{4}, 0, \ldots)$, and $u_3 = (0, 0, 0, 0, \frac{1}{8}, 0 \ldots)$. Using u_n, construct the sequence $z_n := \sum_{j=1}^{n} u_j$. Then $z_1 = (\frac{1}{2}, 0, \ldots)$, $z_2 = (\frac{1}{2}, 0, \frac{1}{4}, 0, \ldots)$, $z_3 = (\frac{1}{2}, 0, \frac{1}{4}, 0, \frac{1}{8}, 0 \ldots)$, $\ldots$. has the desirable property that it is increasing and none of the z_n are order units. However, this sequence order converges to $\bar{z} = (\frac{1}{2}, 0, \frac{1}{4}, 0, \frac{1}{8}, 0, \frac{1}{16}, \ldots)$, which is not an order unit. At this point, one might be tempted to simply define $u_0 = e - \bar{z}$ and then redefine $\hat{z}_n := u_0 + z_n$. This gives $\hat{z}_1 = (1, 1, \frac{3}{4}, 1, \frac{7}{8}, \ldots)$, $\hat{z}_2 = (1, 1, 1, 1, \frac{7}{8}, \ldots), \ldots$ and $\hat{z}_n \xrightarrow{\circ} e$. Unfortunately, each z_n is an order unit since $u_0 = (\frac{1}{2}, 1, \frac{3}{4}, 1, \frac{7}{8}, \ldots)$ is an order unit and thus the sequence $\{\hat{z}_n\}$ fails the premise of Theorem 4.3.8. More care is necessary.

Constructing the appropriate sequence of non-order units requires two additional steps. First, project e onto the band generated by the sequence $\{z_n\}$. This projection is well defined and is equal to $P_B(e) := \sup\{e \wedge nz_n : n = 1, 2, \ldots\} = (1, 0, 1, 0, 1, \ldots)$. Define a new sequence $w_n := e \wedge nz_n$. Each w_n has the same support as z_n and each of its positive components increase to 1 as $n \to \infty$, so $w_1 = (\frac{1}{2}, 0, \dots)$, $w_2 = (1, 0, \frac{1}{2}, 0, \dots)$, $w_3 = (1, 0, \frac{3}{4}, 0, \frac{3}{8}, 0 \dots)$, $\dots$, and $w_n \overset{\circ}{\to} P_B(e)$. Second, construct the vector $u_0 = e - P_B(e) = (0, 1, 0, 1, 0, \dots)$. The w_n for each n and u_0 are not order units and the sequence $\{u_0 + w_n\}$ is an increasing sequence of non-order units such that $u_0 + w_n \overset{\circ}{\to} e$. ◁

Theorem 4.3.12. *Assume E is an infinite dimensional Riesz space that is either σ-order complete or has the projection property. If E contains an order unit $e \succ \theta_E$, then E contains an increasing sequence x_n of non-order units with $x_n \overset{\circ}{\to} e$.*

Proof. By hypothesis E is σ-order complete or has the projection property. Therefore E is Archimedean by the main inclusion theorem (see Theorem 25.1 of Luxemburg and Zaanen [1971] and our Figure 4.1). Since E is Archimedean and infinite dimensional, E contains an infinite set $\{u_n\}_{n=1}^{\infty}$ of pairwise disjoint elements by Lemma 4.3.10. By the definition of disjoint, we can assume that this sequence is positive. Since the u_n are pairwise disjoint none are order units by Lemma 4.5.4.

Define a new sequence of vectors $\{z_n\}$ from the sequence of vectors $\{u_n\}$ by

$$z_n := \sum_{k=1}^{n} u_k. \tag{4.3.6}$$

Since each element of $\{z_n\}$ is a linear combination of the elements of $\{u_n\}$, Lemma 4.5.5 implies that none of the $\{z_n\}$ are order units. Note that $\theta_E \preceq z_n \uparrow$ since each u_k is a positive vector. The sequence $\{z_n\}$ either converges in order to e or it does not. If $\{z_n\}$ does order converge to e the proof is complete with the desired sequence $\{x_n\} := \{z_n\}$.

Assume $\{z_n\}$ does not order converge to e. Define B to be the band generated by $\{z_n\}$. Since E is either $\sigma-$order complete or has the projection property and $\theta_E \preceq z_n \uparrow$, Theorem 28.3 in Luxemburg and Zaanen [1971] implies that B is a projection band and

$$P_B(e) = \sup\{e \wedge n z_m \,:\, m, n = 1, 2, \dots\} = \sup\{e \wedge n z_n \,:\, n = 1, 2, \dots\}. \tag{4.3.7}$$

Define $w_n := e \wedge nz_n$ for $n \in \mathbb{N}$. Clearly, $\theta_E \preceq \{w_n\} \uparrow$ so (4.3.7) implies $\{w_n\} \uparrow P_B(e) \preceq e$. By definition, $w_n \preceq nz_n$ for each $n \in \mathbb{N}$. Since the z_n are not order units, nz_n are not order units and $\theta_E \preceq w_n \preceq nz_n$ implies that none of the $\{w_n\}$ are order units. If $e = P_B(e)$, the proof is complete with $\{x_n\} := \{w_n\}$, since $\{w_n\}$ is an increasing sequence of non-order units that order converges to e.

Assume $e \neq P_B(e)$. Since $P_B(e) \preceq e$ it follows that $P_B(e) \prec e$ and $u_0 := e - P_B(e) \succ \theta_E$. Consider the sequence $\{u_0 + w_n\}$. Since $\{w_n\} \uparrow P_B(e)$, it follows from the definition of u_0 that $\{u_0 + w_n\} \uparrow e$. Because B is a projection band, u_0 is disjoint from the $z_1, z_2, \ldots,$ that generate B. Then by Lemma 4.5.5, $u_0 + nz_n$ is not an order unit for all n. Finally, observe that

$$\theta_E \prec u_0 + w_n = u_0 + (e \wedge nz_n) = (u_0 + e) \wedge (u_0 + nz_n) \preceq u_0 + nz_n.$$

Since the $u_0 + nz_n$ are not order units, the $u_0 + w_n$ are not order units. Hence the increasing sequence $\{x_n\} := u_0 + w_n$ of non-order units converges to e. This establishes the result. □

Theorem 4.3.13. *All infinite dimensional Riesz spaces with order units that are either σ-order complete or have the projection property have non-zero, positive, singular dual functionals.*

Proof. Let E be an infinite dimensional Riesz space with an order unit $e \succ \theta_E$ that is either σ-order complete or has the projection property. By Theorem 4.3.12 there exists an increasing sequence $\{x_n\}$ of non-order units such that $x_n \overset{\circ}{\to} e$. Then by Theorem 4.3.8 there exists a non-zero, positive, singular dual functional on E. □

Remark 4.3.14.

We now relate our results to the main inclusion theorem of Luxemburg and Zaanen expressed in Figure 4.1. In the proof of Theorem 4.3.12 we used the fact that E was Archimedean to generate an infinite set of pairwise disjoint vectors. Unfortunately the Archimedean property, as shown in Example 4.3.5, is not sufficient to guarantee the existence of singular dual functionals. If one moves to the left in Figure 4.1 to either $\sigma-$order complete or the projection property, then by Theorem 4.3.13 either one of these conditions is sufficient to guarantee singular functionals in the presence of an order unit. However, as illustrated in Example 4.3.9, these conditions are not necessary to guarantee the existence of singular functionals.

The principal projection property lies between $\sigma-$order complete/projection property and Archimedean. It is possible that the principal projection property is sufficient to guarantee singular functionals in the presence of an order unit, but this remains an open question. We do know that the principal projection property is not a necessary condition for singular functionals. Indeed, $C[0,1]$ is an infinite dimensional Riesz space with order unit that does not satisfy the principal projection property (see Example (v) on page 140 of Luxemburg and Zaanen [1971]) but has singular dual functionals (see page 147 of Zaanen [1997]). $\triangleleft$

Remark 4.3.15.

Theorem 4.3.13 uses an order-algebraic approach to show that a large class of infinite dimensional Riesz spaces satisfy Theorem 4.3.8 and therefore have singular functionals in their order duals. A similar, but less general result, using topological methods is given in Appendix 4.5.4. There we show that if E is a σ-order complete Banach lattice with an order unit then $E^{\sim}$ admits singular functionals. However, this result requires the additional structure of being a Banach lattice. This is not required in Theorem 4.3.13. By Theorem 9.28 of Aliprantis and Border [2006], any order complete Riesz space E can be equipped with a norm that defines a Banach lattice. However, unlike Theorem 4.3.13, this requires order completeness as opposed to σ-order completeness. Moreover, Theorem 4.3.13 also handles the case where the space satisfies the projection property only and is possibly not even σ-order complete. We also reiterate that Theorem 4.3.8 applies to all infinite dimensional Riesz spaces, providing a general condition for the existence of singular dual functionals. $\triangleleft$

4.4 Resolving the Slater Conundrum

Connecting our algebraic constraint qualification based on core points to the existence of singular dual functionals reveals a significant modeling issue in infinite dimensional programming. We show the tradeoff between the sufficiency of interior point conditions for zero duality gap and the difficulties of working with singular dual functionals. The question thus remains how to "get around" the Slater conundrum and still work in spaces like ℓ_∞ and L_∞, which are otherwise desirable for modeling purposes.

We provide sufficient conditions to bypass the conundrum for general infinite dimensional linear programs over Riesz spaces with structure (4.3.1). Consider

$$\begin{aligned} \inf_{x} \quad & \langle x, \varphi \rangle \\ \text{s.t.} \quad & A(x) \succeq_{Y^+} b \\ & x \succeq_{X^+} \theta_X \end{aligned} \tag{RLP}$$

where $A : X \to Y$ is a linear map and X and Y are both infinite-dimensional Riesz spaces where $\varphi \in X^\sim$, $b \in Y$, Y^+ is a pointed convex convex cone with nonempty core in Y, and X^+ is a pointed convex cone in X. When the underlying vector space is clear, we drop the subscripts on the orderings $\succeq_{X^+}$ and $\succeq_{Y^+}$ for ease of notation. The order dual of (RLP) is

$$\begin{aligned} \sup_{\psi} \quad & \langle b, \psi \rangle \\ \text{s.t.} \quad & A^\sim(\psi) \preceq_{X^+} \varphi \\ & \psi \in (Y^+)^\sim \end{aligned} \tag{RLP$^\sim$}$$

where $A^\sim : Y^\sim \to X^\sim$ is the *order adjoint* of A defined by $\langle x, A^\sim(\psi)\rangle = \langle A(x), \psi\rangle$ for all $x \in X$ and $\psi \in Y^\sim$, and $(Y^+)^\sim$ is the order dual cone of Y^+.

Remark 4.4.1.

The linear program (RLP) requires that $x \in X^+$, a condition not included of previous formulations in this chapter. However, this condition is easily relaxed via a standard argument in linear programming. $\triangleleft$

Our investigation is motivated by the following question. Given an instance of (RLP) where there exists an optimal dual solution to (RLP$^\sim$), does there always exist an optimal dual solution that has no singular component? If we can affirmatively answer this, then we say that we have *resolved* the Slater conundrum. One sufficient condition for resolving the conundrum is:

$$\text{If } \psi^* = \psi_c^* + \psi_s^* \text{ is optimal to (RLP}^\sim\text{) with } \psi_c^* \in Y_c^\sim, \psi_s^* \in Y_s^\sim \text{ then } \psi_c^* \text{ is also optimal.} \tag{4.4.1}$$

In this section we provide two sets of sufficient conditions that guarantee (4.4.1) holds and thus resolve the Slater conundrum. The first set of conditions is inspired by Theorem 5.1 of Ponstein [1981] for problems in ℓ_∞.

Theorem 4.4.2. *Consider an instance of* (RLP) *where A and b are such that the following two conditions hold: for all positive singular linear functionals $\psi_s \in (Y^+)_s^\sim$*

$$A^\sim(\psi_s) \succeq \theta \tag{4.4.2}$$

and

$$\langle b, \psi_s \rangle \leq 0. \tag{4.4.3}$$

Then condition (4.4.1) *holds.*

Proof. Let $\psi^* = \psi_c^* + \psi_s^*$ be an optimal solution to (RLP$^\sim$). The proof proceeds in two steps. We first show that ψ_c^* is itself a feasible dual solution using (4.4.2). Second, we show that $\langle b, \psi_c^* \rangle \geq \langle b, \psi^* \rangle$ using (4.4.3). This establishes that ψ_c^* is also an optimal solution and so (4.4.1) holds. To establish the first step note that $\varphi \succeq A^{\sim}(\psi^*) = A^{\sim}(\psi_c^*) + A^{\sim}(\psi_s^*) \succeq A^{\sim}(\psi_c)$ where the first inequality follows from the feasibility of ψ^*, the equality follows from the linearity of the order adjoint and the second inequality follows by (4.4.2). To establish the second step note that $\langle b, \psi^* \rangle = \langle b, \psi_c^* \rangle + \langle b, \psi_s^* \rangle \leq \langle b, \psi_c^* \rangle$ where the equality follows from linearity and the inequality follows from (4.4.3). □

Ponstein [1981] considers problems where $Y = \ell_\infty$ and uses the notion *singular nonpositive* to condition A and b so that (4.4.2) and (4.4.3) hold. A sequence $y = (y_1, y_2, \dots) \in \ell_\infty$ is singularly nonpositive if every accumulation point of y is nonpositive. Lemma 4.2 of Ponstein [1981] shows that if $\psi \in (\ell_\infty^{\sim})_s$ then $\psi(y) \leq 0$ for all singularly nonpositive $y \in \ell_\infty$. Ponstein considers a linear program of the form (RLP) where $A : X \to \ell_\infty$ is represented by a doubly infinite matrix (also denoted A) and X is a sequence space that contains the unit vectors e^i for $i = 1, 2, \dots$. Theorem 5.1 of Ponstein [1981] then shows that if $-b$ and the columns of A (themselves sequences) are singularly nonpositive then (4.4.2) and (4.4.3) hold, thus establishing (4.4.1) via our Theorem 4.4.2.

A sufficient condition to guarantee that the columns of A are singularly nonpositive is to require that that the matrix A has finitely many nonzero entries in each column. This sufficient condition was used by Ghate and Smith [2013]; Romeijn et al. [1992]; Romeijn and Smith [1998] in their work on the duality of countably infinite programs. By using this sufficient condition, they were able to ignore the contributions of singular dual functionals.

What can be done to resolve the conundrum when either (4.4.2) or (4.4.3) fail? We provide one further set of sufficient conditions, based on a decomposition of the dual problem, that draws partial inspiration from Shapiro's approach to the duality of conic optimization problems over L_∞ in Shapiro [2005]. The idea is to use the orthogonality of $Y_c^{\sim}$ and $Y_s^{\sim}$ and assume a particular structure on A that allows the dual problem to be decomposed into a "continuous" subproblem and a "singular" subproblem.

Theorem 4.4.3. *Assume the primal problem* (RLP) *is feasible with* $\varphi \in X_c^\sim$ *and the order adjoint* $A^\sim$ *satisfies* $A^\sim((Y^+)_n^\sim) \subseteq X_n^\sim$ *and* $A^\sim((Y^+)_s^\sim) \subseteq X_s^\sim$*; that is, the order adjoint maps positive* σ*-order continuous linear functionals to* σ*-order continuous linear functionals and positive singular linear functionals to singular linear functionals. Then condition* (4.4.1) *holds.*

Proof. Let $\psi^* = \psi_c^* + \psi_s^*$ be an optimal solution to (RLP$^\sim$). The proof proceeds by decomposing (RLP$^\sim$) into two subproblems, one involving only ψ_c^* and one involving only ψ_s^*. First, we decompose the constraint $\psi^* = \psi_c^* + \psi_s^* \succeq \theta$ in (RLP$^\sim$). Since $\psi_c^* \perp \psi_s^*$, Theorem 4.5.6 implies $\psi_c^* \succeq \theta$ and $\psi_s^* \succeq \theta$. Of course, $\psi_c^* \succeq \theta$ and $\psi_s^* \succeq \theta$ implies $\psi_c^* + \psi_s^* \succeq \theta$, and so these two conditions are equivalent. Second, we decompose the constraint $A^\sim(\psi^*) \preceq \varphi$. Note that $A^\sim(\psi^*) \preceq \varphi$ holds if and only if $\varphi - A^\sim(\psi_c^*) - A^\sim(\psi_s^*) \succeq \theta$. By assumption $\varphi \in X_c^\sim$ and $A^\sim(\psi_c^*) \in X_c^\sim$ since $\psi_c^* \in (Y^+)_c^\sim$ so $\varphi - A^\sim(\psi_c^*) \in X_c^\sim$. Moreover, by assumption $A^\sim(\psi_s^*) \in X_s^\sim$ since $\psi_s^* \in (Y^+)_s^\sim$. Then by Theorem 4.5.6, $\varphi - A^\sim(\psi_c^*) - A^\sim(\psi_s^*) \succeq \theta$ implies $\varphi - A^\sim(\psi_c^*) \succeq \theta$ and $-A^\sim(\psi_s^*) \succeq \theta$.

The above decompositions imply that ψ^* is an optimal solution to (RLP$^\sim$) if and only if ψ_c^* is an optimal solution to

$$\begin{aligned} \sup_{\psi} \quad & \langle b, \psi_c \rangle \\ \text{s.t.} \quad & A^\sim(\psi_c) \preceq \varphi \\ & \psi_c \in (Y^+)_c^\sim \end{aligned} \tag{RLP$_c^\sim$}$$

and ψ_s^* is an optimal solution to

$$\begin{aligned} \sup_{\psi} \quad & \langle b, \psi_s \rangle \\ \text{s.t.} \quad & A^\sim(\psi_s) \preceq \theta \\ & \psi_s \in (Y^+)_s^\sim \end{aligned} \tag{RLP$_s^\sim$}$$

with $v(\text{RLP}^{\sim}) = v(\text{RLP}^{\sim}_{c}) + v(\text{RLP}^{\sim}_{s})$.

The feasible region to $(\text{RLP}^{\sim}_{s})$ is a cone since if ψ_s is feasible then $\lambda\psi_s$ is feasible for all $\lambda \geq 0$. This follows directly from the linearity of the adjoint $A^{\sim}(\lambda\psi_s) = \lambda A^{\sim}(\psi_s) \preceq \theta$ for all $\lambda \geq 0$. Therefore, if there exists a feasible solution ψ_s to $(\text{RLP}^{\sim}_{s})$ such that $\langle b, \psi_s \rangle > 0$ then $(\text{RLP}^{\sim}_{s})$ is unbounded. But (RLP) is feasible, so $\infty > v(\text{RLP}) \geq v(\text{RLP}^{\sim})$ by weak duality. Thus $\langle b, \psi^{*}_{s} \rangle = v(\text{RLP}^{\sim}_{s}) = 0$ which implies $\psi_s = \theta$ is also optimal solution to $(\text{RLP}^{\sim}_{s})$. Then $v(\text{RLP}^{\sim}) = v(\text{RLP}^{\sim}_{c})$ and $\psi^{*}_{c} + \theta = \psi^{*}_{c}$ is also an optimal solution to $(\text{RLP}^{\sim})$. This establishes (4.4.1). □

Remark 4.4.4.

The approaches in Theorem 4.4.2 and Theorem 4.4.3 are different and not implied by each other. Theorem 4.4.3 puts no condition on the right-hand-side b and so is more general in this regard. Theorem 4.4.2 does not restrict where $A^{\sim}$ maps continuous or singular linear functionals, and is thus more general in this direction. We leave for future work the implications of Theorem 4.4.3 for particular optimization problems in spaces such as ℓ_∞ or L_∞. ◁

4.5 Appendices

4.5.1 The Convex Core Topology

In this appendix we show Corollary 4.2.2 follows from Theorem 4.2.1 under a special topology that captures the inherent algebraic structure of the vector space.

Let Y be a vector space. A subset A of Y is *algebraically open* if $\text{cor}(A) = A$. Let $\mathcal{B}$ denote the set of all algebraically open *convex* sets that contain the origin θ_X. The set $\mathcal{B}$ forms the neighborhood basis of the origin for the topology $\tau(\mathcal{B})$ where

$$U \in \tau(\mathcal{B}) \text{ if and only if } \forall y \in U,\ \exists B \in \mathcal{B} \text{ s.t. } y + B \subseteq U. \tag{4.5.1}$$

Following Day [1973] we call $\tau(\mathcal{B})$ the *convex core topology* of Y. It has also been called the *natural topology* by Klee [1963] and generates the locally convex topological vector space $(Y, \tau(\mathcal{B}))$.

Since $(Y, \tau(\mathcal{B}))$ is a locally convex topological vector space we can apply Theorem 4.2.1 to a problem with $(Y, \tau(\mathcal{B}))$ as a constraint space. Corollary 4.2.2 is then a direct consequence of the following proposition, whose proof is straightforward and thus omitted.

Proposition 4.5.1. *Let Y be a vector space with positive cone P and $\tau(\mathcal{B})$ the natural topology on Y. Then $\operatorname{int}(P) = \operatorname{cor}(P)$ and the topological dual Y^* under the natural topology is the algebraic dual Y'.*

4.5.2 Comparison with quasi-relative interior

Borwein and Lewis [1992] propose the quasi-relative interior as a useful generalization of the notion of interior for constraint qualifications of convex programs. Let X be a topological vector space with topological dual X^* and let $A \subseteq X$ be a convex set. The quasi-relative interior of a set A is $\operatorname{qri}(A) := \{x \in A : \operatorname{cl}(\operatorname{cone}(A - x)) \text{ is a linear subspace of } X\}$.

By Corollary 4.2.2 (CP) has zero duality gap with its Lagrangian dual if there exists an $x \in X$ such that $-G(x) \in \operatorname{cor}(P)$. Example 4.5.2 below shows that when the condition $-G(x) \in \operatorname{cor}(P)$ is replaced with $-G(x) \in \operatorname{qri}(P)$, then there may be a positive duality gap. Thus Corollary 4.2.2 is not immediately subsumed by results in the literature on constraint qualifications involving the quasi-relative interior. See Example 21.1 in Boţ [2010] for an example where the existence of a quasi-relative interior point is not sufficient to guarantee a zero duality gap with the Fenchel dual.

Example 4.5.2.

Consider the semi-infinite linear program

$$
\begin{aligned}
&\min x_1 \\
&-(1/n)x_1 - (1/n)^2 x_2 + (1/n) \leq 0, \quad n = 1, 2, 3, \ldots
\end{aligned}
\tag{4.5.2}
$$

with constraint space $\ell_2(\mathbb{N})$. An optimal primal solution is $(x_1^*, x_2^*) = (1, 0)$ for an optimal primal value of 1. A feasible dual solution is a positive dual functional ψ that satisfies

$$\psi\left(\{-1/n\}_{n=1}^{\infty}\right) = 1 \tag{4.5.3}$$

$$\psi\left(\left\{-(1/n)^2\right\}_{n=1}^{\infty}\right) = 0. \tag{4.5.4}$$

We claim that the algebraic dual is infeasible. By Theorem 4.2.4 and Theorem 4.5.9 every positive dual functional ψ corresponds to an element $\{\psi_n\}_{n=1}^{\infty}$ in $\ell_2(\mathbb{N})$ with $\psi_n \geq 0$ for all n. Thus, (4.5.3) and (4.5.4) amount to $\sum_{n=1}^{\infty} -(1/n)\psi_n = 1$ and $\sum_{n=1}^{\infty} -(1/n)^2\psi_n = 0$. The above inequalities cannot be satisfied by any positive ψ. The first equality requires at least one of the ψ_n to be strictly positive, which violates the second inequality. We conclude that the algebraic dual of (4.5.2) is infeasible.

Since the primal is feasible and dual is infeasible, (4.5.2) has an infinite duality gap. However, there exists an $\bar{x}$ such that $-G(\bar{x}) \in \text{qri}(P)$ where G is the linear map defining the constraints of (4.5.2) and $P = (\ell_2(\mathbb{N}))^+$. Let $\bar{x} = (2, 0)$. Then $\bar{y} = -G(\bar{x})$ is the sequence $\{1/n\}_{n=1}^{\infty}$. Example 3.11(i) of Borwein and Lewis [1992] shows that $\text{qri}(P) = \{y \in \ell_2(\mathbb{N}) : y_n > 0 \text{ for all } n\}$. Hence $\bar{y} \in \text{qri}(P)$. $\triangleleft$

4.5.3 Properties of Riesz Spaces

This appendix contains results about Riesz spaces used in the chapter that are not readily found in other references.

Proposition 4.5.3. *Assume E is Riesz space with positive cone E^+. If $x, y, z \in E^+$, then*

$x \perp (y+z) \Rightarrow x \perp y$ and $x \perp z$.

Proof. By definition $x \perp (y+z)$ implies that $|x| \wedge |y+z| = \theta_E$. Then $x, y, z \in E^+$ implies that $\theta_E = |x| \wedge |y+z| = x \wedge (y+z)$. Therefore $\theta_E = x \wedge (y+z) \succeq x \wedge y \succeq \theta_E$ and $\theta_E = x \wedge (y+z) \succeq x \wedge z \succeq \theta_E$, which implies that $x \perp y$ and $x \perp z$. □

Lemma 4.5.4. Let x and y be elements of the Riesz space E. If $y \neq \theta_E$ and $x \wedge y = \theta_E$, then x is not an order unit.

Proof. Prove the contrapositive and assume x is an order unit. Then there exists a $\lambda > 0$ such that $y \preceq \lambda x$. Then $\lambda x \wedge y = y$ and $y \neq \theta_E$ by hypothesis. By Theorem 8.1(ii) in Zaanen [1997], this implies that $x \wedge y \neq \theta_E$, i.e. x and y are not disjoint. □

Lemma 4.5.5. Let $\{u_n\}_{n=1}^{\infty}$ be a sequence of nonzero disjoint vectors in Riesz space E. No linear combination of $\{u_n\}_{n=1}^{\infty}$ is an order unit in E.

Proof. Since the u_n for $n \in \mathbb{N}$ are nonzero disjoint vectors, it follows from Lemma 4.5.4 that the u_n are not order units. Let $\overline{u} = \sum_{n=1}^{N} \alpha_n u_n$ be an arbitrary linear combination of the elements of $\{u_n\}$. Then Theorems 8.1(ii)-(iii) in Zaanen [1997] imply that $\overline{u} \perp u_{N+1}$. If $\overline{u}$ is an order unit, then there exists a $\lambda > 0$ such that $u_{N+1} \preceq \lambda \overline{u}$. Thus, $\lambda \overline{u}$ and u_{N+1} are not disjoint since $u_{N+1} \wedge \lambda \overline{u} = u_{N+1} \neq \theta$. Then by Theorem 8.1(ii) in Zaanen [1997], $\overline{u}$ and u_{N+1} are not disjoint, yielding a contradiction. Since $\overline{u}$ was an arbitrary linear combination of the the elements of $\{u_n\}$, no linear combination of the elements of the u_n is an order unit. □

Theorem 4.5.6. *(Cone Decomposition) If E is a Riesz space with $y, z \in E$ such that $y \perp z$, then $y + z \succeq \theta$ if and only if $y \succeq \theta$ and $z \succeq \theta$*

Proof. ($\Leftarrow$) If $y \succeq \theta$ and $z \succeq \theta$ then $y + z \succeq \theta$. ($\Rightarrow$) If $y + z \succeq \theta$, then

$$y + z = |y + z| = |y - z|$$

where the second equality follows from Theorem 8.12 (2) of Aliprantis and Border [2006] and the fact that y and z are disjoint. Then by Theorem 8.6 (8) of Aliprantis and Border [2006],

$$y \wedge z = \frac{1}{2}(y + z - |y - z|)$$

but $y + z = |y + z| = |y - z|$ gives

$$y \wedge z = \frac{1}{2}(y + z - |y - z|) = \frac{1}{2}(|y - z| - |y - z|) = \theta.$$

If $y \wedge z = \theta$, then $y \succeq (y \wedge z) = \theta$ and $z \succeq (y \wedge z) = \theta$. □

4.5.4 Banach Lattices

Theorem 4.3.13 uses an order-algebraic approach to prove that infinite dimensional, σ-order complete Riesz spaces with order units have order duals that contain singular dual functionals. Alternatively, similar, but weaker results can be proved using topological methods. First some definitions.

Definition (Banach lattice)**.** Let E be a Riesz space with norm $||\cdot||$. We say $||\cdot||$ is a *lattice norm* if for $x, y \in E$ with $|x| \leq |y|$, then $||x|| \leq ||y||$. A Riesz space equipped with a lattice norm is called a *normed Riesz space.* If the lattice norm is norm-complete (that is, Cauchy sequences converge in norm) then E is a *Banach lattice.*

Definition (AM-space)**.** A lattice norm on a Riesz space E is an *M-norm* if $x, y \succeq \theta_E$ implies $||x \vee y|| = \max\{||x||, ||y||\}$. A Banach lattice with an M-norm is an *AM-space.*

Definition (Order Continuous Norm)**.** The lattice norm $\|\cdot\|$ is an order continuous norm if $x_\alpha \downarrow \theta_E$ implies $\|x_\alpha\| \downarrow 0$.

Definition (σ-Order Continuous Norm)**.** The lattice norm $\|\cdot\|$ is a σ- order continuous norm if $x_n \downarrow \theta_E$ implies $\|x_n\| \downarrow 0$.

Theorem 4.5.7 (Aliprantis and Border [2006], Theorem 9.28). *If E is either a Banach lattice or an order complete Riesz space, then for each $x \in E$ the principal ideal E_x, equipped with the norm $||y||_\infty = \inf\{\lambda > 0 : |y| \preceq \lambda |x|\} = \min\{\lambda \geq 0 : |y| \preceq \lambda x\}$ is an AM-space, with order unit $|x|$.*

If E is an order complete Riesz space with an order unit $e \succ \theta_E$, then the principal ideal E_e is all of E. Therefore, Theorem 4.5.7 implies that every order complete Riesz space with order unit is a Banach lattice, indeed an AM-space when equipped with norm topology $|| \cdot ||_\infty$.

Theorem 4.5.8 (Aliprantis and Burkinshaw [2006], Corollary 4.4). *All lattice norms that make a Riesz space a Banach lattice are equivalent.*

Theorem 4.5.9 (Aliprantis and Burkinshaw [2006], Corollary 4.5). *If E is a Banach lattice then $E^* = E^\sim$ where E^* is the norm dual and $E^\sim$ is the order dual.*

Theorem 4.5.10 (Aliprantis and Burkinshaw [2006], Theorem 4.51; or Wnuk [1999], Theorem 1.1). *For an arbitrary σ-order complete Banach lattice E the following statements are equivalent: (i) E does not have order continuous norm and (ii) There exists an order bounded disjoint sequence of E^+ that does not converge in norm to zero.*

Theorem 4.5.11 (Zaanen [1983], Theorem 103.9). *A Banach lattice E has an order continuous norm if and only if E is σ-order complete and E has a σ-order continuous norm.*

Theorem 4.5.12 (Zaanen [1983], Theorem 102.8). *A Banach lattice E has a σ-order continuous norm if and only if $E^* = E_c^\sim$.*

The logic for Theorem 4.5.13 given below is inspired by page 48 of Wnuk [1999].

Theorem 4.5.13. *If an infinite dimensional vector space E is either a σ-order complete Banach lattice or an order complete Riesz space and E contains an order unit $e \succ \theta_E$, then there exists a nonzero singular functional in the algebraic dual E', the order dual $E^\sim$ and the norm dual E^*.*

Proof. By Theorem 4.5.7, $(E, ||\cdot||_\infty)$ is an AM-space with order unit where $||\cdot||_\infty$ is defined by

$$\|x\|_\infty := \inf\{\lambda > 0 \, : \, |x| \preceq \lambda e\}. \tag{4.5.5}$$

By Theorem 4.5.8 all lattice norms that make E a Banach lattice are equivalent, and so without loss, take E to be the Banach lattice $E = (E, \|\cdot\|_\infty)$.

By hypothesis, E is infinite dimensional. Then by Lemma 4.3.10, E contains an infinite sequence $\{x_n\}_{n=1}^{\infty}$ of pairwise disjoint elements. By definition of disjointness, we can assume this sequence is positive. By Theorem 8.1(ii) in Zaanen [1997] the sequence defined as

$$y_n := \frac{x_n}{\|x_n\|_\infty} \tag{4.5.6}$$

is still positive pairwise disjoint. By the definition of y_n in (4.5.6), $\|y_n\|_\infty = 1$ which implies by definition of the $\|\cdot\|_\infty$ in (4.5.5) that $|y_n| \preceq e$. Thus the sequence $\{y_n\}_{n=1}^{\infty}$ is order bounded by e. Hence $\{y_n\}_{n=1}^{\infty}$ is an order bounded sequence of positive pairwise disjoint elements that norm converges to 1. Therefore condition (ii) of Theorem 4.5.10 holds and this implies E does not have an order continuous norm.

Since E is σ-order complete without an order continuous norm, Theorem 4.5.11 implies that E does not have a σ-continuous norm. Therefore, Theorem 4.5.12 implies $E^* \neq E_c^{\sim}$ and there exists a nonzero singular dual functional in the norm dual of E. Since $E^* = E^{\sim} \subseteq E'$, the norm dual, the order dual and the algebraic dual of E contain a nonzero singular dual functional. □

REFERENCES

D. Abreu, D. Pearce, and E. Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica: Journal of the Econometric Society*, pages 1041–1063, 1990.

D. Adelman and D. Klabjan. Duality and existence of optimal policies in generalized joint replenishment. *Mathematics of Operations Research*, 30(1):28–50, 2005.

P. Aghion, M.P. Espinosa, and B. Jullien. Dynamic duopoly with learning through market experimentation. *Economic Theory*, 3(3):517–539, 1993.

C.D. Aliprantis and K.C. Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, second edition, 2006.

C.D. Aliprantis and D.J. Brown. Equilibria in markets with a Riesz space of commodities. *Journal of Mathematical Economics*, 11(2):189–207, 1983.

C.D. Aliprantis and O. Burkinshaw. *Positive Operators*, volume 119. Springer, 2006.

C.D. Aliprantis and R Tourky. *Cones and Duality*. American Mathematical Society, second edition, 2007.

E.J. Anderson and P. Nash. *Linear Programming in Infinite-Dimensional Spaces: Theory and Applications*. Wiley, 1987.

V.F. Araman and R. Caldentey. Dynamic pricing for nonperishable products with demand learning. *Operations research*, 57(5):1169–1188, 2009.

Y. Aviv and A. Pazgal. A partially observed markov decision process for dynamic pricing. *Management Science*, 51(9):1400–1416, 2005.

A. Basu, K. Martin, and C. T. Ryan. On the sufficiency of finite support duals in semi-infinite linear programming. *Operations Research Letters*, 42(1):16–20, 2014.

J.P. Benoit and V. Krishna. Finitely repeated games. *Econometrica: Journal of the Econometric Society*, pages 905–922, 1985.

D. Bertsimas and G. Perakis. *Dynamic pricing: A learning approach*. Springer, 2006.

O. Besbes and A. Zeevi. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research*, 57(6):1407–1420, 2009.

O. Besbes and A. Zeevi. On the minimax complexity of pricing in a changing environment. *Operations research*, 59(1):66–79, 2011.

Gabriel Bitran and René Caldentey. An overview of pricing models for revenue management. *Manufacturing & Service Operations Management*, 5(3):203–229, 2003.

L.E. Blume and D. Easley. What has the rational learning literature taught us? In A. Kirman and M. Salmon, editors, *Learning and rationality in economics*. B. Blackwell, 1995.

J.M. Borwein and A.S. Lewis. Partially finite convex programming, Part I: Quasi relative interiors and duality theory. *Mathematical Programming*, 57(1-3):15–48, 1992.

R.I Boţ. *Conjugate Duality in Convex Optimization*, volume 637. Springer, 2010.

J. Broder and P. Rusmevichientong. Dynamic pricing under a general parametric choice model. *Operations Research*, 60(4):965–980, 2012.

J.G. Carlsson. Dividing a territory among several vehicles. *INFORMS Journal on Computing*, 24(4):565–577, 2012.

A.X. Carvalho and M.L. Puterman. Learning and pricing in an internet environment with binomial demands. *Journal of Revenue and Pricing Management*, 3(4):320–336, 2005.

W.L. Cooper, T. Homem-de Mello, and A.J. Kleywegt. Learning and pricing with models that do not explicitly incorporate competition. *Operations research*, 63(1):86–103, 2015.

M.M. Day. *Normed linear spaces*. Springer, 1973.

G. Debreu. Valuation equilibrium and Pareto optimum. *Proceedings of the National Academy of Sciences of the United States of America*, 40(7):588, 1954.

M.A.H. Dempster. On stochastic programming II: Dynamic problems under risk? *Stochastics: An International Journal of Probability and Stochastic Processes*, 25(1):15–42, 1988.

A.V. den Boer. Dynamic pricing with multiple products and partially specified demand distribution. *Mathematics of operations research*, 39(3):863–888, 2014.

A.V. den Boer. Dynamic pricing and learning: historical origins, current research, and new directions. *Surveys in operations research and management science*, 20(1):1–18, 2015.

A.V. den Boer and B. Zwart. Simultaneously learning and optimizing using controlled variance pricing. *Management Science*, 60(3):770–783, 2013.

R.J. Duffin and L.A. Karlovitz. An infinite linear program with a duality gap. *Management Science*, 12(1):122–134, 1965.

V.F. Farias and B. Van Roy. Dynamic pricing with a prior on market response. *Operations Research*, 58(1):16–29, 2010.

K. Fong, O. Gossner, J. Horner, and Y. Sannikov. Efficiency in a repeated prisoners' dilemma with imperfect private monitoring, 2008.

D.P. Foster and H.P. Young. Learning, hypothesis testing, and Nash equilibrium. *Games and Economic Behavior*, 45(1):73–96, 2003.

D.P. Foster and H.P. Young. Regret testing: learning to play Nash equilibrium without knowing you have an opponent. *Theoretical Economics*, 1(3):341–367, 2006.

D. Fudenberg and D.K. Levine. *The theory of learning in games*, volume 2. MIT press, 1998.

G. Gallego and G. Van Ryzin. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management science*, 40(8):999–1020, 1994.

A. Ghate. *Markov Chains, Game Theory, and Infinite Programming: Three Paradigms for Optimization of Complex Systems*. University of Michigan (thesis), 2006.

A. Ghate and R.L. Smith. A linear programming approach to non stationary infinite-horizon Markov decision processes. *Operations Research*, 61:413–425, 2013.

A. Grad. Quasi-relative interior-type constraint qualifications ensuring strong Lagrange duality for optimization problems with cone and affine constraints. *Journal of Mathematical Analysis and Applications*, 361(1):86–95, 2010.

E.J. Green and R.H. Porter. Noncooperative collusion under imperfect price information. *Econometrica: Journal of the Econometric Society*, pages 87–100, 1984.

R.C. Grinold. *Continuous Programming*. PhD thesis, University of California, Berkeley, 1968.

J.E. Harrington, Jr. Experimentation and learning in a differentiated-products duopoly. *Journal of Economic Theory*, 66(1):275–288, 1995.

J.M. Harrison, N.B. Keskin, and A. Zeevi. Bayesian dynamic pricing policies: Learning and earning under a binary prior distribution. *Management Science*, 58(3):570–586, 2012.

S. Hart, A. Mas-Colell, et al. Stochastic uncoupled dynamics and Nash equilibrium. *Games and Economic Behavior*, 57(2):286–303, 2006.

R.B. Holmes. *Geometric Functional Analysis and its Applications*. Springer, 1975.

B. Janjgava and S. Slobodyan. *Duopoly Competition, Escape Dynamics and Non-cooperative Collusion*. Economics Institute, Academy of Sciences of the Czech Republic, 2011.

J.S. Jordan. Bayesian learning in normal form games. *Games and Economic Behavior*, 3(1): 60–81, 1991.

E. Kalai and E. Lehrer. Rational learning leads to Nash equilibrium. *Econometrica: Journal of the Econometric Society*, pages 1019–1045, 1993.

D.F. Karney. Duality gaps in semi-infinite linear programming – an approximation problem. *Mathematical Programming*, 20(1):129–143, 1981.

G. Keller and S. Rady. Price dispersion and learning in a dynamic differentiated-goods duopoly. *RAND Journal of Economics*, pages 138–165, 2003.

N. Keskin and A. Zeevi. Dynamic pricing with an unknown demand model: Asymptotically optimal semi-myopic policies. *Operations Research*, 62:1142–1167, 2014.

A. Kirman and M. Salmon. *Learning and rationality in economics.* B. Blackwell, 1995.

V. Klee. The Euler characteristic in combinatorial geometry. *American Mathematical Monthly*, 70:119–127, 1963.

T.L. Lai and H. Robbins. Iterated least squares in multiperiod control. *Advances in Applied Mathematics*, 3(1):50–73, 1982.

M.S. Lobo and S. Boyd. Pricing and learning with uncertain demand. In *INFORMS Revenue Management Conference*, 2003.

D.G. Luenberger. *Optimization by Vector Space Methods*. Wiley, 1969.

W.A.J. Luxemburg and A.C. Zaanen. *Riesz Spaces I.* North-Holland, 1971.

A.M. Manelli and D.R. Vincent. Optimal procurement mechanisms. *Econometrica: Journal of the Econometric Society*, pages 591–620, 1995.

C. Marlats. A folk theorem for stochastic games with finite horizon. *Economic Theory*, 58 (3):485–507, 2015.

A. McLennan. Price dispersion and incomplete learning in the long run. *Journal of Economic dynamics and control*, 7(3):331–347, 1984.

J.P. Ponstein. On the use of purely finitely additive multipliers in mathematical programming. *Journal of Optimization Theory and Applications*, 33(1):37–55, 1981.

J.P. Ponstein. *Approaches to the Theory of Optimization*, volume 77. Cambridge University Press, 2004.

M.C. Pullan. An algorithm for a class of continuous linear programs. *SIAM Journal on Control and Optimization*, 31:1558, 1993.

D. Rahman. The dilemma of the cypress and the oak tree. Technical report, Discussion paper, 2014.

R.T. Rockafellar. *Conjugate Duality and Optimization*. SIAM, Philadelphia, PA, 1974.

R.T. Rockafellar and R.J.B. Wets. Stochastic convex programming: singular multipliers and extended duality, singular multipliers and duality. *Pacific J. Math*, 62(2):507–522, 1976a.

R.T. Rockafellar and R.J.B. Wets. Stochastic convex programming: relatively complete recourse and induced feasibility. *SIAM Journal on Control and Optimization*, 14(3):574–589, 1976b.

H. E. Romeijn, R.L. Smith, and J.C. Bean. Duality in infinite dimensional linear programming. *Mathematical programming*, 53(1-3):79–97, 1992.

H.E. Romeijn and R.L. Smith. Shadow prices in infinite-dimensional linear programming. *Mathematics of operations research*, 23(1):239–256, 1998.

M. Rothschild. A two-armed bandit theory of market pricing. *Journal of Economic Theory*, 9(2):185–202, 1974.

B.C. Schipper. Strategic teaching and learning in games. *Available at SSRN 2594193*, 2015.

A. Shapiro. On the duality theory of conic linear programs. In M. Goberna and M. Lopez, editors, *Semi-Infinite Programming: Recent Advances*, pages 135–165. Springer, 2005.

C. Simon. *Dynamic pricing with demand learning under competition.* PhD thesis, Massachusetts Institute of Technology, 2007.

G.J. Stigler. A theory of oligopoly. *The Journal of Political Economy*, pages 44–61, 1964.

J.A. Tropp et al. Freedman's inequality for matrix martingales. *Electron. Commun. Probab*, 16:262–270, 2011.

N. Williams. *Escape dynamics in learning models.* PhD thesis, University of Chicago, Dept. of Economics, 2001.

W. Wnuk. *Banach Lattices with Order Continuous Norms.* Polish Scientific Publishers, 1999.

Y. Yamamoto. Stochastic games with hidden states, second version. 2015.

Yuichi Yamamoto. Individual learning and cooperation in noisy repeated games. *The Review of Economic Studies*, 81(1):473–500, 2014.

H.P. Young. *Strategic Learning and Its Limits*, volume 2002. Oxford University Press on Demand, 2004.

A.C. Zaanen. *Riesz Spaces II.* North-Holland, 1983.

A.C. Zaanen. *Introduction To Operator Theory in Riesz Spaces.* Springer, 1997.

M. Zinkevich, M. Johanson, M. Bowling, and C. Piccione. Regret minimization in games with incomplete information. In *Advances in neural information processing systems*, pages 1729–1736, 2007.