

Distributionally robust inventory control when demand is a martingale

Linwei Xin

University of Illinois at Urbana-Champaign, lxin@illinois.edu, <https://sites.google.com/site/linweilawrencexinshomepage2/>

David A. Goldberg

Georgia Institute of Technology, dgoldberg9@isye.gatech.edu, <http://www2.isye.gatech.edu/~dgoldberg9/>

Demand forecasting plays an important role in many inventory control problems. To mitigate the potential harms of model misspecification in this context, various forms of distributionally robust optimization have been applied. Although many of these methodologies suffer from the problem of time-inconsistency, the work of Klabjan, Simchi-Levi and Song [83] established a general time-consistent framework for such problems by connecting to the literature on robust Markov decision processes.

Motivated by the fact that many forecasting models exhibit very special structure, as well as a desire to understand the impact of positing different dependency structures, in this paper we formulate and solve a time-consistent distributionally robust multi-stage newsvendor model which naturally unifies and robustifies several inventory models with demand forecasting. In particular, many simple models of demand forecasting have the feature that demand evolves as a martingale (i.e. expected demand tomorrow equals realized demand today). We consider a robust variant of such models, in which the sequence of future demands may be any martingale with given mean and support. Under such a model, past realizations of demand are naturally incorporated into the structure of the uncertainty set going forwards.

We explicitly compute the minimax optimal policy (and worst-case distribution) in closed form, by combining ideas from convex analysis, probability, and dynamic programming. We prove that at optimality the worst-case demand distribution corresponds to the setting in which inventory may become obsolete at a random time, a scenario of practical interest. To gain further insight, we prove weak convergence (as the time horizon grows large) to a simple and intuitive process. We also compare to the analogous setting in which demand is independent across periods (analyzed previously in Shapiro [115]), and identify interesting differences between these models, in the spirit of the price of correlations studied in Agrawal et al. [2].

Key words: inventory control, distributionally robust optimization, martingale, dynamic programming, robust Markov decision process, demand forecasting

1. Introduction

The fundamental problem of managing an inventory over time in the presence of stochastic demand is one of the core problems of Operations Research, and related models have been studied since at least the seminal work of Edgeworth [45]. In many practical settings of interest, demands are correlated over time (cf. Scarf [111, 112], Iglehart and Karlin [75]). As a result, there is a vast literature investigating inventory models with correlated demand, including: studies of the so-called bull-whip effect (cf. Ryan [109], Chen et al. [32], Lee et al. [85]); models with Markov-modulated demand (cf. Feldman [49], Iglehart and Karlin [75], Kalyon [81]); and models with forecasting, including models in which demand follows an auto-regressive / moving average (ARMA) or exponentially smoothed process (cf. Johnson and Thompson [80], Lovejoy [87], Blinder [27], Pindyck [100], Miller [93], Badinelli [6], Gallego and Ozer [55], Levi et al. [86], Chao [38]); and models obeying the Martingale Model of Forecast Evolution (MMFE) and its many generalizations (cf. Heath and Jackson [69], Graves et al. [63], Toktay and Wein [122], Erkip et al. [47], Iida and Zipkin [76], Lu, Song and Regan [88], Milner and Kouvelis [94]). Although several of these works offer insights into the qualitative impact of correlations on the optimal policy (and associated costs) when managing an

inventory over time, these results are typically proven under very particular distributional assumptions, which assume perfect knowledge of all relevant distributions. This is potentially a significant problem, since various authors have previously noted that model misspecification when demand is correlated can lead to very sub-optimal policies (cf. Badinelli [6]). Indeed, the use of such time series and forecasting models in Operations Research practice is well-documented, and concerns over the practical impact of model misspecification have been raised repeatedly in the forecasting and Operations Research literature (cf. Fildes [50], Fildes et al. [51]).

One approach taken in the literature to correcting for such model uncertainty is so-called distributionally robust optimization. In this framework, one assumes that the joint distribution (over time) of the sequence of future demands belongs to some set of joint distributions, and solves the minimax problem of computing the control policy which is optimal against a worst-case distribution belonging to this set. Such a distributionally robust approach is motivated by the reality that perfect knowledge of the exact distribution of a given random process is rarely available (cf. Dupacová [43, 44], Prékopa [103], Žáčková [131]). A typical minimax formulation is as below:

$$\min_{x \in \mathcal{X}} \max_{Q \in \mathfrak{M}} \mathbb{E}_Q [f(x, \xi)],$$

where \mathcal{X} is a set of decisions and \mathfrak{M} is a set of probability measures. The objective is to pick a decision x that minimizes the average cost of f under a worst-case distribution.

The application of such distributionally robust approaches to the class of inventory control problems was pioneered in Scarf [110], where it was assumed that \mathfrak{M} contains all probability measures whose associated distributions satisfy certain moment constraints. Such an approach has been taken to many variants of the single-stage model since then (cf. Gallego and Moon [57], Gallego [54], Perakis and Roels [99], Gallego and Moon [58], Yue et al. [132], Hanasusanto et al. [66], Zhu et al. [134]), and the single-stage distributionally robust model is quite broadly understood.

The analogous questions become more subtle in the multi-stage setting, due to questions regarding the specification of uncertainty in the underlying joint distribution. There have been two approaches taken in the literature, depending on whether the underlying optimization model is *static* or *dynamic* in nature. In a static formulation, one specifies a family of joint distributions for demand over time, typically by e.g. fixing various moments and supports, or positing that the distribution belongs to some appropriately defined ball, and then solves an associated global minimax optimization problem (cf. Delage and Ye [40], See and Sim [114], Popescu [102], Bertsimas et al. [17], Dupacová [43, 44], Bertsimas, Gupta and Kallus [19]). Such static formulations generally cannot be decomposed and solved by dynamic programming (DP), because the distributional constraints do not contain sufficient information about how the distribution behaves under conditioning. Put another way, such models generally do not allow for the incorporation of realized demand information into model uncertainty / robustness going forwards (i.e. re-optimization in real time), and are generally referred to as time-inconsistent in the literature (cf. Ahmed, Cakmak and Shapiro [4], Iancu, Petrik and Subramanian [72], Xin, Goldberg and Shapiro [130], Epstein and Schneider [46]). This inability to incorporate realized demand information may make such approaches undesirable for analyzing models which explicitly consider the forecasting of future demand (cf. Klabjan, Simchi-Levi and Song [83], Iancu, Petrik and Subramanian [72], Shapiro [116]). We note that although some of these static formulations have been able to model settings in which information is revealed over time and/or temporal dependencies manifest through time series (cf. See and Sim [114], Lam [84], Gupta [65]), the fundamental inability to incorporate realized demand (in the sense of time in-consistency) remains.

Alternatively, in a dynamic formulation, the underlying family of potential joint distributions must admit a certain decomposition which ensures a desirable behavior under conditioning, and a resolution by DP. The existence of such a decomposition is generally referred to as the *rectangularity*

property (cf. Iyengar [78], Iancu, Petrik and Subramanian [72], Epstein and Schneider [46], Shapiro [118]). For example, the set of all joint distributions for the vector of demands (D_1, D_2) such that (s.t.): $\mathbb{E}[D_1] = 1$, $\mathbb{E}[D_2|D_1] = D_1$, and (D_1, D_2) has support on the non-negative integers is rectangular, since the feasible set of joint distributions may be decomposed as follows. To each possible realized value d for D_1 (i.e. each non-negative integer), we may associate a fixed collection $S(d)$ of possible conditional distributions for D_2 (i.e. those distributions with mean d and support on the non-negative integers). Furthermore, every feasible distribution for the vector (D_1, D_2) may be constructed by first selecting a feasible distribution \mathcal{D} for D_1 (i.e. any distribution with mean one and support on the non-negative integers), and for each element d in the support of \mathcal{D} , setting the conditional distribution of D_2 (given $\{D_1 = d\}$) to be any fixed element of $S(d)$. Moreover, the set of joint distributions constructible in this manner is precisely the set of feasible distributions for the vector (D_1, D_2) . Alternatively, if we had instead required that $\mathbb{E}[D_1] = \mathbb{E}[D_2] = 1$, and $\mathbb{E}[D_1 D_2] = 2$ (with the same support constraint), the corresponding set of feasible joint distributions would *not* be rectangular, as it may be verified that such a decomposition is no longer possible. For a formal definition and more complete / precise description of the rectangularity property, we refer the reader to cf. Iyengar [78], Iancu, Petrik and Subramanian [72], Epstein and Schneider [46], Xin, Goldberg and Shapiro [130], Nilim and El Ghaoui [96], Bielecki et al. [13], Shapiro [118], and note that since various communities have studied several closely related notions at different times, a complete consensus on a common rigorous definition has not yet been reached. As discussed in Nilim and El Ghaoui [96], we note that such dynamic formulations are closely related to the theory of stochastic games (cf. Altman et al. [5]), and have also been studied within the mathematical finance community (cf. Bayraktar et al. [9], Bayraktar and Yao [10], Glasserman and Xu [60], Cheng and Riedel [34], Föllmer and Leukert [52], Nishimura and Ozaki [97], Riedel [106], Trojanowska and Kort [123]).

There are several works which formulate DP approaches to distributionally robust inventory models (cf. Gallego [56], Ahmed, Cakmak and Shapiro [4], Choi and Ruszczyński [35], Shapiro [115], Xin, Goldberg and Shapiro [130]). More generally, such dynamic problems can typically be formulated as so-called robust Markov decision processes (MDP) (cf. Iyengar [78], Nilim and El Ghaoui [96], Wiesemann, Kuhn and Rustem [129]), akin to robust formulations considered previously within the control community (cf. Hansen and Sargent [67]). To the best of our knowledge, the only such work which considers applications to correlated demands and forecasting models within Operations Research is that of Klabjan, Simchi-Levi and Song [83]. In that work, the authors make the connection between dynamic distributionally robust inventory models and robust MDP, in a very general sense, and prove e.g. the optimality of state-dependent (s, S) policies, where the state incorporates the entire history of demand. The authors also consider a more specialized model in which the demand in period t may come from any distribution within a certain ball of the empirical histogram of all past demands. We note that many simple forecasting models common in the Operations Research and statistics literature (e.g. linear time series) have considerable special structure, and to the best of our knowledge precisely how this structure would manifest within the general framework of Klabjan, Simchi-Levi and Song [83] (and associated general framework of robust MDP) remains an open question.

On a related note, to the best of our knowledge, there seems to have been no systematic study of the impact of positing different joint dependency structures in such multi-stage distributionally robust inventory control problems, e.g. comparing the properties and cost of a minimax optimal policy. The quest to develop such an understanding in the broader context of stochastic optimization (not specifically inventory control) was initiated in Agrawal et al. [2], where the authors define the so-called *price of correlations* as the ratio between the optimal minimax value when all associated random variables (r.v.) are independent, and the setting where they may take any joint distribution belonging to an allowed family. Although the authors do not look specifically at any

inventory problems, they stress the general importance of understanding how positing different joint distribution uncertainty impacts the underlying stochastic optimization. Also, although several authors have demonstrated that the family of non-stationary history-dependent s-S policies is optimal for a broad class of minimax inventory problems (cf. Ahmed, Cakmak and Shapiro [4], Notzon [98], Klabjan, Simchi-Levi and Song [83]), the impact of the posited dependency structure on the corresponding s-S policy remains poorly understood. We note that although these questions remain open, the general theory of sensitivity analysis for (distributionally robust) optimization problems is very well developed (cf. Bonans and Shapiro [28]), and that different yet related questions are the subject of several recent investigations (cf. Gupta [65]).

Combining the above, we are led to the following question.

QUESTION 1. Can we construct effective dynamic distributionally robust variants of the structured time series and forecasting models common in the Operations Research and statistics literature? Furthermore, can we develop a theory of how positing different dependency structures impacts the optimal policy and cost for such models?

Highly relevant to any discussion of distributionally robust optimization (either static or dynamic) is the vast literature on classical (i.e. deterministic) robust optimization, in which the only constraints made are on the supports of the associated r.v.s (cf. Ben-Tal, El Ghaoui and Nemirovski [14], Bertsimas, Brown and Caramanis [18]). In this setting, the worst-case distribution is always degenerate, namely a point-mass on a particular worst-case trajectory. Such models often lead to tractable global optimization problems for fairly complex models, and have been successfully applied to several inventory settings (cf. Kasugai and Kasegai [82], Aharon et al. [3], Ben-Tal et al. [16], Bertsimas, Iancu and Parrilo [21], Bertsimas and Thiele [24], Ben-Tal et al. [15], Carrizosaa et al. [31]). Indeed, these models tend to be more tractable than corresponding distributionally robust formulations, where the question of tractability (under both static and dynamic formulations) may be more sensitive to the particular modeling assumptions (cf. See and Sim [114], Bertsimas, Sim, and Zhang [23], Wiesemann, Kuhn and Rustem [129], Delage and Ye [40]). In spite of their potential computational advantages, one drawback of such classical robust approaches is that they may be overly conservative, and unable to capture the stochastic nature of many real-world problems (cf. See and Sim [114]). We note that the precise relationship between classical robust and distributionally robust approaches remains an intriguing open question. On a related note, questions of robust time series models and their applications to inventory models, similar in spirit to those we will consider in the present work, were recently studied in Carrizosaa et al. [31], at least for single-stage models. There the authors formulated a general notion of robust time series, and prove that the corresponding single-stage minimax newsvendor problem can be formulated as a convex optimization problem. As regards connections to our earlier discussion of static and dynamic formulations in the distributionally robust setting, we note that the potential modeling benefits of being able to dynamically update one's uncertainty sets has also been considered in the deterministic robust setting (cf. Bertsimas and Vayanos [25]).

Also relevant to questions involving distributionally robust optimization is the search for “optimal” probability bounds. In particular, some of the most celebrated results in probability theory are of the form “any random vector with property A has probability at most p of belonging to set S” for appropriate choices of A, p , S. That obtaining tight versions of these bounds can be phrased as an appropriate distributionally robust optimization problem is by now well-known within the Operations Research community, especially since this connection was clarified in the seminal work of Bertsimas and Popescu [22]. Such problems, e.g. those considered in Bertsimas and Popescu [22], typically involve formulating a static optimization problem in which one specifies knowledge or bounds on certain moments, covariances, supports, etc., and frames finding a “worst-case” distribution satisfying these conditions as a convex optimization problem. Recent work of Beigblock and

Nutz [11] has brought to light the fact that certain probability inequalities, specifically well-known inequalities for martingales such as the celebrated Burkholder inequality, can be similarly formulated using distributionally robust optimization, but are more amenable to a dynamic formulation, due to the fundamentally conditional structure of the martingale property. Another related line of work is that on so-called optimal transport problems, especially that subset focusing on so-called optimal martingale transport, and we refer the interested reader to Dolinsky and Soner [41] and Biegblock et al. [12] for details and additional references.

We close this overview by noting that there is a vast literature on time series in the statistics community (cf. Box, Jenkins, and Reinsel [29]), and we make no attempt to survey that literature here. These investigations include considerable work on robust approaches (cf. Maronna et al. [90]), where different lines of research use different notions of robustness (cf. Stockinger and Dutter [121]). Most relevant to our own investigations are the many such works related to distributional misspecification, which has been investigated within the statistics community for over 50 years (cf. Tukey [125], Huber [71], Stockinger and Dutter [121], Medina and Ronchetti [92]), and includes the analysis of minimax notions of robustness (cf. Franke [53], Luz and Moklyachuk [89]) and so-called robust Kalman filters (cf. Ruckdeschel [108]). In spite of this vast literature, it is worth noting that even recently, there have been several calls for the development of more tools and methodologies for the robust analysis of time series data. For example, in J. De Gooijer and R. Hyndman [61], it is noted that a persistent problem in the application of time series models is “**the frequent misuse of methods based on linear models with Gaussian i.i.d. distributed errors.**” A similar commentary was given very recently in Guerrier et al. [64], in which the authors note that “**the robust estimation of time series parameters is still a widely open topic in statistics for various reasons.**” Such calls for the development of new methodologies further serve to underscore the potential impact of combining ideas from robust optimization with time series analysis, e.g. through Question 1.

1.1. Our contributions

In this paper we take a step towards answering Question 1 by formulating and solving a dynamic distributionally robust multi-stage newsvendor model which naturally unifies the analysis of several inventory models with demand forecasting. In particular, we consider a dynamic distributionally robust multi-stage inventory control problem (with backlogged demand) in which the sequence of future demands may be any martingale with given mean and support (assumed the same in every period). Our contributions are four-fold. First, we explicitly compute the minimax optimal policy (and associated worst-case distribution) in closed form. Our main proof technique involves a non-trivial induction, combining ideas from convex analysis, probability, and DP. Second, we prove that at optimality the worst-case demand distribution corresponds to the setting in which the inventory may become obsolete at a random time, a scenario of practical interest which has been studied previously in the literature (cf. Pierskalla [101], Brown [30], Joglekar and Lee [79], Song and Zipkin [120], Cerag [37]). Third, to gain further insight into our explicit solution, we prove weak convergence (as the time horizon grows large) to a simple and intuitive process. Fourth, in the spirit of Agrawal et al. [2], we compare to the analogous setting in which demand is independent across periods (analyzed previously in Shapiro [115]), and identify interesting differences between these two models. In particular, we prove that when the initial inventory level is zero, the minimax optimal cost is maximized under the dependency structure of the independent-demand model. Furthermore, we compute the limiting ratio (as the time horizon grows large) of the minimax optimal value under the martingale-demand model to that under the independent-demand model, which has a simple closed form and evaluates to exactly $\frac{1}{2}$ in the perfectly symmetric case. Interestingly, we find that under different initial conditions, the situation may in fact be reversed, showing that such comparative statements may be quite subtle.

1.2. Outline of paper

Our paper proceeds as follows. We begin by reviewing the model with independent demand (analyzed in Shapiro [115]) in Section 2.1. We then introduce our distributionally robust model with martingale dependency structure in Section 2.2, and prove the validity of a certain DP formulation (very similar to the robust MDP framework of Iyengar [78]). We present our main results in Section 2.3. The proof of our main result, i.e. explicit solution to our distributionally robust martingale model, is given in Section 3. In Section 4 we present our results regarding asymptotic analysis and weak convergence. We provide a summary of our results, directions for future research, and concluding remarks in Section 5. We also provide a technical appendix, which contains the proofs of several results used throughout the paper, in Section 6.

2. Main results

Before stating our main results, we briefly review the inventory model analyzed in Shapiro [115], which we refer to as the independent-demand model.

2.1. Independent-demand model

Consider the following distributionally robust inventory control problem with backlogging, finite time horizon T , strictly positive linear backorder per-unit cost b , and a per-unit holding cost of 1, where we note that assuming a holding cost equal to 1 is without loss of generality (w.l.o.g.) due to a simple scaling argument. Let D_t be the demand in period t , and x_t be the inventory level at period t after placing a non-negative order, $t \in [1, T]$, where we note that the order must be placed at the start of period t before D_t is known. For a t -dimensional vector \mathbf{x} (potentially of real numbers, r.v.s, functions, etc.), let $\mathbf{x}_{[s]} = (x_1, \dots, x_s)$, and $\mathbf{x}_{[t_1, t_2]} = (x_{t_1}, \dots, x_{t_2})$. In that case, we require that x_t is a measurable function of $D_{[t-1]}$, $x_t \geq x_{t-1} - D_{t-1}$ for $t \in [2, T]$, and that x_1 is a real constant, where $x_0 \in [0, U]$ denotes the initial inventory level and $x_1 \geq x_0$.

For $t \in [1, T]$, the cost incurred in period t equals

$$C_t \triangleq b[D_t - x_t]_+ + [x_t - D_t]_+,$$

where $z_+ \triangleq \max(z, 0)$, $z_- \triangleq \max(-z, 0)$, $z \in \mathbb{R}$. As a notational convenience, we also define $C(x, d) \triangleq b(d - x)_+ + (x - d)_+$.

As in Iyengar [78], to avoid later possible concerns about measurability, throughout we make the minor technical assumption that all relevant demand distributions have rational support, and all policies π considered are restricted to order rational quantities (i.e. not irrational quantities such as $\sqrt{2}$), and the relevant problem primitives x_0, μ, U, b (corresponding to initial inventory level, mean demand, bound on support of demand, and backorder penalty respectively) also belong to \mathcal{Q}^+ , with $x_0, \mu \in [0, U]$. Under this assumption, it may be easily verified (due to the finiteness and non-negativity of all relevant cost functions) that all relevant functions arising in the various DP which we shall later define are sufficiently regular, e.g. measurable, integrable, and finite-valued. For clarity of exposition, and in light of this assumption, we do not dwell further on related questions of measurability, and instead refer the interested reader to the excellent reference Puterman [104] for further discussion.

We now formally and more precisely define the relevant notion of admissible policy in our setting. As our analysis will be restricted to the setting in which there is a uniform upper bound on the support of demand (denoted U), for simplicity we will incorporate U directly into several of our definitions. Let $[0, U]_{\mathcal{Q}} \triangleq [0, U] \cap \mathcal{Q}$. Then we define an admissible policy π to be a T -dimensional vector of functions $\{x_t^\pi, t \in [1, T]\}$, s.t. x_1^π is a constant satisfying $x_1^\pi \geq x_0^\pi \triangleq y_1^\pi \triangleq x_0$; and for $t \in [2, T]$, x_t^π is a deterministic map from $[0, U]_{\mathcal{Q}}^{t-1}$ to \mathcal{Q} , s.t. for $t \in [1, T-1]$ and $\mathbf{d} = (d_1, \dots, d_t) \in [0, U]_{\mathcal{Q}}^t$, $x_{t+1}^\pi(\mathbf{d}) \geq y_{t+1}^\pi(\mathbf{d}) \triangleq x_t^\pi(\mathbf{d}_{[t-1]}) - d_t$, which is equivalent to requiring that a non-negative amount of

inventory is ordered in each period. Sometimes as a notational convenience, we will denote x_1^π by $x_1^\pi(\mathbf{d}_{[0]})$, or more generally as $x_1^\pi(q)$ for $q \in [0, U]_{\mathcal{Q}}$.

Let $\hat{\Pi}$ denote the family of all such admissible policies. Note that once a particular policy $\pi \in \hat{\Pi}$ is specified, the associated costs $\{C_t^\pi, t \in [1, T]\}$ are explicit functions of the vector of demands, and we sometimes make this dependence explicit with the notation $C_t^\pi(D_{[t]})$, where $C_t^\pi(D_{[t]}) = C(x_t^\pi(D_{[t-1]}), D_t)$.

Since we will exclusively consider settings in which the demand in every period belongs to $[0, U]$, and the initial inventory belongs to $[0, U]$, it will be convenient to prove that in such a setting one may w.l.o.g. restrict to policies which always order up to a level in $[0, U]$. In particular, let Π denote that subset of $\hat{\Pi}$ consisting of policies π s.t. $x_t^\pi(\mathbf{d}_{[t-1]}) \in [0, U]$ for all $t \in [1, T]$ and $\mathbf{d} \in [0, U]_{\mathcal{Q}}^T$ whenever $x_0 \in [0, U]_{\mathcal{Q}}$ (which again is to be assumed throughout). Then the following lemma provides the desired result, where we defer the relevant proof to the Technical Appendix in Section 6.

LEMMA 1. *For every $\hat{\pi} \in \hat{\Pi}$, there exists $\pi \in \Pi$ s.t. for all $\mathbf{d} = (d_1, \dots, d_T) \in [0, U]_{\mathcal{Q}}^T$ and $t \in [1, T]$, $C_t^\pi(\mathbf{d}_{[t]}) \leq C_t^{\hat{\pi}}(\mathbf{d}_{[t]})$.*

Let us say that a probability measure Q on an appropriate measurable space is rationally supported if it places probability 1 on t -dimensional rational vectors for some finite t , where we note that all measures which we consider will have this property. For such a measure Q , let $\text{supp}(Q)$ (i.e. the support of Q) denote the (possibly infinite) set of rational vectors which Q assigns strictly positive probability, and D^Q denote a r.v. (or vector if $\text{supp}(Q)$ is not one-dimensional) distributed as Q . For $\mu \in [0, U]_{\mathcal{Q}}$, let $\mathfrak{M}(\mu)$ be the collection of all rationally supported probability measures Q s.t. $\text{supp}(Q) \subseteq [0, U]_{\mathcal{Q}}$, and $E[D^Q] = \mu$. Furthermore, let \mathbf{IND}_T denote the collection of all T -dimensional product measures s.t. all T marginal distributions belong to $\mathfrak{M}(\mu)$. When there is no ambiguity, we will suppress the dependence on T , simply writing \mathbf{IND} . In words, the joint distribution of demand belongs to \mathbf{IND} iff the demand is independent across time periods, and the demand in each period has support a subset of $[0, U]_{\mathcal{Q}}$ and mean μ . Then the particular optimization problem considered in Shapiro [115] Example 5.1.2 is $\inf_{\pi \in \hat{\Pi}} \sup_{Q \in \mathbf{IND}} \mathbb{E}_Q[\sum_{t=1}^T C_t^\pi]$, which by Lemma 1 is equivalent to

$$\inf_{\pi \in \Pi} \sup_{Q \in \mathbf{IND}} \mathbb{E}_Q[\sum_{t=1}^T C_t^\pi]. \quad (2.1)$$

We note that although the exact problem considered in Shapiro [115] does not explicitly restrict the family of policies and demand distributions to \mathcal{Q} , it may be easily verified that the same results go through in this setting (under our assumptions on the rationality of the relevant problem parameters). In Example 5.1.2 of Shapiro [115], the author proves that due to certain structural properties, Problem (2.1) can be reformulated as a DP, as we now describe. We note that although the family of measures \mathbf{IND} does not possess the so-called rectangularity property, here the decomposition follows from the fact that convexity and independence imply that irregardless of the inventory level and past realizations of demand, a worst-case demand distribution always has support $\{0, U\}$, where the precise distribution is then dictated by the requirement of having expectation μ . To formally define the relevant DP, we now define a sequence of functions $\{V^t, t \geq 1\}, \{f^t, t \geq 1\}, \{g^t, t \geq 1\}$. In general, such a DP would be phrased in terms of the so-called ‘‘cost-to-go’’ functions, with the t -th such function representing the remaining cost incurred by an optimal policy during periods t through T , subject to the given state at time t . Here and throughout, we use the fact that the backorder and holding costs are the same in every period to simplify the relevant concepts and notations. In particular, due to this symmetry and the associated self-reducibility which it induces, for all $T \geq 1$, it will suffice to define the aforementioned cost-to-go function for the first time period only. Indeed, in the following function definitions, $V^T(y, \mu)$ will coincide with the optimal value of Problem (2.1) when the initial inventory level is y and the associated mean

is μ (we leave the dependence on U implicit), where f^T, g^T have analogous interpretations. We note that to prevent later potential confusion and ambiguity regarding the order of operations such as supremum and expectation, we explicitly write out certain expectations as relevant sums over appropriate domains. Then the relevant DP is as follows.

$$\begin{aligned} f^1(x_1, d_1) &\triangleq C(x_1, d_1) & , & \quad g^1(x_1, \mu) \triangleq \sup_{Q_1 \in \mathfrak{M}(\mu)} \sum_{q_1 \in [0, U]_{\mathcal{Q}}} f^1(x_1, q_1) Q_1(q_1), \\ V^1(y_1, \mu) &\triangleq \inf_{\substack{z_1 \geq y_1 \\ z_1 \in [0, U]_{\mathcal{Q}}}} g^1(z_1, \mu) & , & \quad Q^1(x_1, \mu) \triangleq \arg \max_{Q_1 \in \mathfrak{M}(\mu)} \sum_{q_1 \in [0, U]_{\mathcal{Q}}} f^1(x_1, q_1) Q_1(q_1), \\ \Phi^1(y_1, \mu) &\triangleq \arg \min_{\substack{z_1 \geq y_1 \\ z_1 \in [0, U]_{\mathcal{Q}}}} g^1(z_1, \mu); \end{aligned} \tag{2.2}$$

and for $T \geq 2$,

$$\begin{aligned} f^T(x_T, d_T) &\triangleq C(x_T, d_T) + V^{T-1}(x_T - d_T, \mu) & , & \quad g^T(x_T, \mu) \triangleq \sup_{Q_T \in \mathfrak{M}(\mu)} \sum_{q_T \in [0, U]_{\mathcal{Q}}} f^T(x_T, q_T) Q_T(q_T), \\ V^T(y_T, \mu) &\triangleq \inf_{\substack{z_T \geq y_T \\ z_T \in [0, U]_{\mathcal{Q}}}} g^T(z_T, \mu) & , & \quad Q^T(x_T, \mu) \triangleq \arg \max_{Q_T \in \mathfrak{M}(\mu)} \sum_{q_T \in [0, U]_{\mathcal{Q}}} f^T(x_T, q_T) Q_T(q_T), \\ \Phi^T(y_T, \mu) &\triangleq \arg \min_{\substack{z_T \geq y_T \\ z_T \in [0, U]_{\mathcal{Q}}}} g^T(z_T, \mu). \end{aligned} \tag{2.3}$$

We have given the argument of the relevant functions with superscript T the subscript T (e.g. x_T, d_T are the arguments for f^T) to prevent any possible ambiguities when expanding the relevant recursions, although this is of course not strictly necessary. We do note that conceptually, if the length of the time horizon is fixed to T , then the functions with superscript T (e.g. V^T) actually correspond to the cost-to-go in period 1 (not period T), and the corresponding arguments (e.g. x_T, d_T) correspond to the inventory level and demand in period 1. More generally, for $s \in [0, T-1]$, V^{T-s} will correspond to the cost-to-go in period $s+1$, and the corresponding arguments (e.g. x_{T-s}, d_{T-s}) correspond to the inventory level and demand in period $s+1$. We have chosen our notational scheme to maximize clarity of exposition, and again note that several of the resulting simplifications do rely on the symmetries / time-homogeneities of the specific problem considered.

To illustrate how the relevant functions evolve, note that $V^2(y_2, \mu)$ equals

$$\inf_{\substack{z_2 \geq y_2 \\ z_2 \in [0, U]_{\mathcal{Q}}}} \sup_{Q_2 \in \mathfrak{M}(\mu)} \sum_{q_2 \in [0, U]_{\mathcal{Q}}} (C(z_2, q_2) + V^1(z_2 - q_2, \mu)) Q_2(q_2),$$

which itself equals

$$\inf_{\substack{z_2 \geq y_2 \\ z_2 \in [0, U]_{\mathcal{Q}}}} \sup_{Q_2 \in \mathfrak{M}(\mu)} \sum_{q_2 \in [0, U]_{\mathcal{Q}}} \left(C(z_2, q_2) + \inf_{\substack{z_1 \geq z_2 - q_2 \\ z_1 \in [0, U]_{\mathcal{Q}}}} \sup_{Q_1 \in \mathfrak{M}(\mu)} \sum_{q_1 \in [0, U]_{\mathcal{Q}}} C(z_1, q_1) Q_1(q_1) \right) Q_2(q_2).$$

We now review the results of Lemma 5.1 and Example 5.1.2 of Shapiro [115], which characterizes the optimal policy, value, and worst-case distribution (against an optimal policy) for Problem 2.1, and relates the solution to DP formulation (2.2) - (2.3). Recall that a policy $\pi \in \Pi$ is said to be a

base-stock policy if there exist constants $\{x_t^*, t \in [1, T]\}$, s.t. $x_t^\pi = \max\{y_t^\pi, x_t^*\}$ for all $t \in [1, T]$. Let us define

$$\chi_{\text{IND}}(\mu, U, b) \triangleq \begin{cases} 0 & \text{if } \mu \leq \frac{U}{b+1}, \\ U & \text{if } \mu > \frac{U}{b+1}; \end{cases}$$

$$\text{Opt}_{\text{IND}}^T(\mu, U, b) \triangleq \begin{cases} Tb\mu & \text{if } \mu \leq \frac{U}{b+1}, \\ T(U - \mu) & \text{if } \mu > \frac{U}{b+1}. \end{cases}$$

Let $\mathcal{D}_\mu \in \mathfrak{M}(\mu)$ be the probability measure s.t. $\mathcal{D}_\mu(0) = 1 - \frac{\mu}{U}, \mathcal{D}_\mu(U) = \frac{\mu}{U}$. Then the following is proven in Shapiro [115].

THEOREM 1 ([115]). *For all strictly positive $U, b \in \mathcal{Q}^+, T \geq 1$, and $\mu, x_0 \in [0, U]_{\mathcal{Q}}$, Problem 2.1 always has an optimal base-stock policy π^* , in which the associated base-stock constants $\{x_t^*, t \in [1, T]\}$ satisfy $x_t^* = \chi_{\text{IND}}(\mu, U, b)$ for all $t \in [1, T]$. If $x_0 \leq \chi_{\text{IND}}(\mu, U, b)$, the optimal value of Problem 2.1 equals $\text{Opt}_{\text{IND}}^T(\mu, U, b)$; and for all $x_0 \in [0, U]$, the product measure s.t. all marginals are distributed according to law \mathcal{D}_μ belongs to $\arg \max_{Q \in \text{IND}} \mathbb{E}_Q[\sum_{t=1}^T C_t^{\pi^*}]$. Furthermore, the DP formulation (2.2) - (2.3) can be used to compute these optimal policies, distributions, and values. In particular, for all $x_0, x, \mu \in [0, U]_{\mathcal{Q}}$, $V^T(x_0, \mu)$ is the optimal value of Problem 2.1, $\mathcal{D}_\mu \in \mathcal{Q}^T(x, \mu)$, and $\max(x, \chi_{\text{IND}}(\mu, U, b)) \in \Phi^T(x, \mu)$.*

2.2. Martingale-demand model

In this subsection, we formally define our distributionally robust martingale-demand model, and establish some preliminary results. Let $T, b, \Pi, C_t^\pi, \mathfrak{M}(\mu)$ be exactly as defined for the independent-demand model in Subsection 2.1, i.e. the time horizon, backorder and holding costs, set of admissible policies, cost incurred in period t under policy π , and collection of all probability measures with support $[0, U]_{\mathcal{Q}}$ and mean μ , respectively. Furthermore, let \mathbf{MAR}_T denote the collection of all rationally supported probability measures corresponding to discrete-time martingale sequences (D_1, \dots, D_T) s.t. all t , the marginal distribution of D_t belongs to $\mathfrak{M}(\mu)$. In words, the joint distribution of demand belongs to \mathbf{MAR}_T iff the sequence of demands is a martingale, and the demand in each period has support a subset of $[0, U]_{\mathcal{Q}}$ and mean μ . When there is no ambiguity, we will suppress the dependence on T , simply writing \mathbf{MAR} . Note that by the martingale property, the *conditional* distribution of D_t , conditioned on $D_{[t-1]}$, belongs to $\mathfrak{M}(D_{t-1})$. Then in analogy with Problem 2.1, the optimization problem of interest is $\inf_{\pi \in \hat{\Pi}} \sup_{Q \in \mathbf{MAR}} \mathbb{E}_Q[\sum_{t=1}^T C_t^\pi]$, which by Lemma 1 is equivalent to

$$\inf_{\pi \in \hat{\Pi}} \sup_{Q \in \mathbf{MAR}} \mathbb{E}_Q[\sum_{t=1}^T C_t^\pi]. \quad (2.4)$$

In analogy with the DP formulation (2.2) - (2.3) given for the independent-demand model, we now define an analogous formulation for the martingale-demand setting, and will later prove that this DP formulation indeed corresponds (in an appropriate sense) to Problem 2.4.

$$\begin{aligned} \hat{f}^1(x_1, d_1) &\triangleq C(x_1, d_1) & , & \quad \hat{g}^1(x_1, \mu_1) \triangleq \sup_{Q_1 \in \mathfrak{M}(\mu_1)} \sum_{q_1 \in [0, U]_{\mathcal{Q}}} \hat{f}^1(x_1, q_1) Q_1(q_1), \\ \hat{V}^1(y_1, \mu_1) &\triangleq \inf_{\substack{z_1 \geq y_1 \\ z_1 \in [0, U]_{\mathcal{Q}}}} \hat{g}^1(z_1, \mu_1) & , & \quad \hat{Q}^1(x_1, \mu_1) \triangleq \arg \max_{Q_1 \in \mathfrak{M}(\mu_1)} \sum_{q_1 \in [0, U]_{\mathcal{Q}}} \hat{f}^1(x_1, q_1) Q_1(q_1), \\ \hat{\Phi}^1(y_1, \mu_1) &\triangleq \arg \min_{\substack{z_1 \geq y_1 \\ z_1 \in [0, U]_{\mathcal{Q}}}} \hat{g}^1(z_1, \mu_1); \end{aligned} \quad (2.5)$$

and for $T \geq 2$,

$$\begin{aligned}
\hat{f}^T(x_T, d_T) &\triangleq C(x_T, d_T) + \hat{V}^{T-1}(x_T - d_T, d_T) & , & \quad \hat{g}^T(x_T, \mu_T) \triangleq \sup_{Q_T \in \mathfrak{M}(\mu_T)} \sum_{q_T \in [0, U]_{\mathcal{Q}}} \hat{f}^T(x_T, q_T) Q_T(q_T), \\
\hat{V}^T(y_T, \mu_T) &\triangleq \inf_{\substack{z_T \geq y_T \\ z_T \in [0, U]_{\mathcal{Q}}}} \hat{g}^T(z_T, \mu_T) & , & \quad \hat{Q}^T(x_T, \mu_T) \triangleq \arg \max_{Q_T \in \mathfrak{M}(\mu_T)} \sum_{q_T \in [0, U]_{\mathcal{Q}}} \hat{f}^T(x_T, q_T) Q_T(q_T), \\
\hat{\Phi}^T(y_T, \mu_T) &\triangleq \arg \min_{\substack{z_T \geq y_T \\ z_T \in [0, U]_{\mathcal{Q}}}} \hat{g}^T(z_T, \mu_T).
\end{aligned} \tag{2.6}$$

To again illustrate how the relevant functions evolve, note that $\hat{V}^2(y_2, \mu_2)$ equals

$$\inf_{\substack{z_2 \geq y_2 \\ z_2 \in [0, U]_{\mathcal{Q}}}} \sup_{Q_2 \in \mathfrak{M}(\mu_2)} \sum_{q_2 \in [0, U]_{\mathcal{Q}}} (C(z_2, q_2) + \hat{V}^1(z_2 - q_2, q_2)) Q_2(q_2),$$

which itself equals

$$\inf_{\substack{z_2 \geq y_2 \\ z_2 \in [0, U]_{\mathcal{Q}}}} \sup_{Q_2 \in \mathfrak{M}(\mu_2)} \sum_{q_2 \in [0, U]_{\mathcal{Q}}} \left(C(z_2, q_2) + \inf_{\substack{z_1 \geq z_2 - q_2 \\ z_1 \in [0, U]_{\mathcal{Q}}}} \sup_{Q_1 \in \mathfrak{M}(q_2)} \sum_{q_1 \in [0, U]_{\mathcal{Q}}} C(z_1, q_1) Q_1(q_1) \right) Q_2(q_2).$$

We now prove that the DP 2.5 - 2.6 indeed corresponds to Problem 2.4, and begin with several additional definitions. Recall that given a probability measure $Q \in \mathbf{MAR}_T$, $D^Q = (D_1^Q, \dots, D_T^Q)$ denotes a random vector distributed as Q . Given $t_1 \leq t_2 \in [1, T]$, we let $Q_{[t_1, t_2]}$ denote the probability measure induced by Q on $(D_{t_1}^Q, \dots, D_{t_2}^Q)$, letting $Q_{[t_2]}$ denote $Q_{[1, t_2]}$, and Q_{t_2} denote $Q_{[t_2, t_2]}$. As a notational convenience, we define Q_0 to be the probability measure which assigns μ probability one. Thus $\text{supp}(Q_0) = \{\mu\}$. Also, we sometimes let μ_0 denote μ . Given two finite vectors $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{y} \in \mathbb{R}^m$, we let $\mathbf{x} : \mathbf{y} \in \mathbb{R}^{n+m}$ denote the concatenation of \mathbf{x} and \mathbf{y} . Namely, $(x : y)_k = x_k$ for $k \in [1, n]$, and $(x : y)_k = y_{k-n}$ for $k \in [n+1, n+m]$. Given $Q \in \mathbf{MAR}_T$, $t \in [1, T]$ and $\mathbf{q} \in \text{supp}(Q_{[t-1]})$, let $Q_{t|\mathbf{q}}$ denote the unique probability measure s.t. for all $d \in [0, U]_{\mathcal{Q}}$, $Q_{t|\mathbf{q}}(d) = \frac{Q_{[t](\mathbf{q}; d)}}{Q_{[t-1](\mathbf{q})}}$. Equivalently, $Q_{t|\mathbf{q}}(d) = \mathbb{P}(D_t^Q = d \mid D_{[t-1]}^Q = \mathbf{q})$. Note that by the martingale property, $Q_{t|\mathbf{q}} \in \mathfrak{M}(q_{t-1})$.

Before precisely stating the desired result, we will need some preliminary results regarding the DP 2.5 - 2.6, which assert the non-emptiness of various sets, as well as a certain axiom-of-choice type result asserting the ability to appropriately select elements from various sets.

LEMMA 2. *For all $t \in [1, T]$, and $x, d \in [0, U]_{\mathcal{Q}}$, $\hat{Q}^t(x, d) \neq \emptyset$, $\hat{\Phi}^t(x, d) \neq \emptyset$. Furthermore, it is possible to assign to every $t \in [1, T]$, and $x, d \in [0, U]_{\mathcal{Q}}$, a measure $\bar{Q}_{x,d}^t \in \hat{Q}^t(x, d)$; and rational number $\bar{\Phi}_{x,d}^t \in \hat{\Phi}^t(x, d)$.*

Although we temporarily defer the proof of Lemma 2, as the lemma will follow from the proof of our main result in which we explicitly construct such $\bar{Q}_{x,d}^t$, $\bar{\Phi}_{x,d}^t$, suppose for now that such an assignment $\bar{Q}_{x,d}^t$, $\bar{\Phi}_{x,d}^t$ has been made in an arbitrary yet fixed manner. We now use these quantities to define a corresponding policy and distribution, which will correspond to the optimal policy, and worst-case distribution at optimality, for Problem 2.4. Let $\hat{\pi}^* = (\hat{x}_1^*, \dots, \hat{x}_T^*) \in \Pi$ denote the following policy, which we define inductively. $\hat{x}_1^* = \bar{\Phi}_{x_0, \mu}^T$. Now, suppose we have completely specified \hat{x}_s^* for all $s \in [1, t]$ for some $t \in [1, T-1]$, and thus have also specified $\hat{y}_s^* \triangleq y_s^{\hat{\pi}^*}$ for all $s \in [1, t+1]$. Then we define \hat{x}_{t+1}^* as follows. For all $\mathbf{d} \in [0, U]_{\mathcal{Q}}^t$, $\hat{x}_{t+1}^*(\mathbf{d}) = \bar{\Phi}_{\hat{y}_{t+1}^*(\mathbf{d}), d_t}^{T-t}$. It may be easily verified from the definition of $\hat{\Phi}$ that this defines a valid policy belonging to Π . Similarly, let $\hat{Q}^* \in \mathbf{MAR}_T$ denote the following probability measure, which we define inductively through its conditional distributions.

$\hat{Q}_1^* = \bar{Q}_{x_0, \mu}^T$. Now, suppose we have completely specified $\hat{Q}_{[t]}^*$. Then we specify $\hat{Q}_{[t+1]}^*$ as follows. For all $\mathbf{d} \in \text{supp}(\hat{Q}_{[t]}^*)$, $\hat{Q}_{t+1|\mathbf{d}}^* = \bar{Q}_{x_{t+1}^{\hat{\pi}^*}(\mathbf{d}), d_t}^{T-t}$. Due to its countable support, it may be easily verified that this uniquely and completely specifies a measure $\hat{Q}^* \in \mathbf{MAR}_T$. As a notational convenience, we let $\hat{\mathbf{D}}^* = (\hat{D}_1^*, \dots, \hat{D}_T^*)$ denote a vector distributed as \hat{Q}^* .

Then we have the following theorem, in analogy with Theorem 1, whose proof we defer to the Technical Appendix in Section 6.

THEOREM 2. *For all strictly positive $U, b \in \mathcal{Q}^+, T \geq 1$, and $\mu, x_0 \in [0, U]_{\mathcal{Q}}$, $\hat{\pi}^*$ is an optimal policy for Problem 2.4, and $\hat{Q}^* \in \arg \max_{Q \in \mathbf{MAR}} \mathbb{E}_Q[\sum_{t=1}^T C_t^{\hat{\pi}^*}]$. Furthermore, $\hat{V}^T(x_0, \mu)$ is the optimal value of Problem 2.4.*

Note that in light of Lemma 2, Theorem 2 implies that computing an optimal policy (and associated worst-case distribution at optimality) for Problem 2.4 is equivalent to solving the DP 2.5 - 2.6, and computing the relevant minimizers and maximizers. Furthermore, in light of Theorem 2, solving Problem 2.4 is thus reduced from a minimax problem to purely a maximization problem (of the expected performance of the policy $\hat{\pi}^*$), over a restricted family of martingales. This puts the problem within the general framework of Beigblock and Nutz [11], in which the authors develop a general methodology, based on concave envelopes, to compute optimal martingale inequalities. Indeed, our resolution of the DP 2.5 - 2.6 can be viewed as an explicit execution of this methodology, which exploits the special structure of $\hat{\pi}^*$. It is also interesting to note that Theorem 2 asserts that one can w.l.o.g. restrict to policies and distributions which are Markovian when one considers the three-dimensional state-space consisting of both the post-ordering inventory level and demand in the previous period, as well as a dimension tracking the current time period, even though in principle both the policy and distribution could have exhibited a more complicated form of history dependence.

2.3. Main results

We now state our main results. We begin by introducing some additional definitions and notations. Whenever possible, we will suppress dependence on the parameters μ, U, b for simplicity. As a notational convenience, let us define all empty products to equal unity, and all empty sums to equal zero. For an event A , let $I(A)$ denote the corresponding indicator. For $j \in [-1, T]$, let

$$A_j^T \triangleq \begin{cases} U \prod_{k=j+1}^{T-1} \frac{k}{b+k} & \text{if } j \leq T-1, \\ \frac{b+T}{T} U & \text{if } j = T; \end{cases}$$

and

$$B_j^T \triangleq \frac{j}{b+T} A_j^T.$$

Note that $A_{-1}^T = B_{-1}^T = B_0^T = 0$, $A_{T-1}^T = B_T^T = U$, and both $\{A_j^T, j \in [-1, T]\}$ and $\{B_j^T, j \in [-1, T]\}$ are increasing in j , and decreasing in T . For $x, \mu \in [0, U]$, let $\mathbf{q}_{x, \mu}^T$ be the unique (up to sets of measure zero) probability measure s.t.

$$\begin{cases} \mathbf{q}_{x, \mu}^T(A_j^T) = \frac{A_{j+1}^T - \mu}{A_{j+1}^T - A_j^T}, \quad \mathbf{q}_{x, \mu}^T(A_{j+1}^T) = \frac{\mu - A_j^T}{A_{j+1}^T - A_j^T} & \text{if } \mu \in (A_j^T, A_{j+1}^T], \quad x \in [0, B_{j+1}^T); \\ \mathbf{q}_{x, \mu}^T(0) = 1 - \frac{\mu}{A_k^T}, \quad \mathbf{q}_{x, \mu}^T(A_k^T) = \frac{\mu}{A_k^T} & \text{if } \mu \in (A_j^T, A_{j+1}^T], \quad x \in [B_k^T, B_{k+1}^T), \quad k \geq j+1; \\ \mathbf{q}_{x, \mu}^T(0) = 1 - \frac{\mu}{U}, \quad \mathbf{q}_{x, \mu}^T(U) = \frac{\mu}{U} & \text{if } \mu = 0 \text{ or } x = U. \end{cases}$$

Also, for $j \in [0, T]$, let

$$G_j^T(x, \mu) \triangleq (T - \frac{b+T}{A_j^T} \mu)x + (T-j)b\mu;$$

in which case we define

$$\Gamma_\mu^T \triangleq \begin{cases} 0 & \text{if } \mu = 0, \\ j+1 & \text{if } \mu \in (A_j^{T+1}, A_{j+1}^{T+1}), j \in [-1, T-1]; \end{cases}$$

and note that $\mu \in (A_{\Gamma_\mu^T - 1}^T, A_{\Gamma_\mu^T}^T)$ for all $T \geq 2$ and $\mu \in (0, U]$, while $A_{\Gamma_0^T - 1}^T = 0$. We also define

$$\chi_{\text{MAR}}^T(\mu, U, b) \triangleq \beta_\mu^T \triangleq B_{\Gamma_\mu^T}^T, \quad \text{Opt}_{\text{MAR}}^T(\mu, U, b) \triangleq G_{\Gamma_\mu^T}^T(\beta_\mu^T, \mu).$$

Then an explicit solution to the DP 2.5 - 2.6, and thus to Problem 2.4, is given as follows.

THEOREM 3. *For all $t \in [1, T]$, and all $x, d \in [0, U]_{\mathcal{Q}}$, $\max(x, \chi_{\text{MAR}}^t(d, U, b)) \in \hat{\Phi}^t(x, d)$, and we may set $\bar{\Phi}_{x,d}^t = \max(x, \chi_{\text{MAR}}^t(d, U, b))$. Thus an optimal policy $\hat{\pi}^*$ for Problem 2.4 is given as follows. $\hat{x}_1^* = \max(x_0, \chi_{\text{MAR}}^T(\mu, U, b))$. For $t \in [1, T-1]$,*

$$\hat{x}_{t+1}^*(\mathbf{d}_{[t]}) = \max(\hat{x}_t^*(\mathbf{d}_{[t-1]}) - d_t, \chi_{\text{MAR}}^{T-t}(d_t, U, b)).$$

Similarly, $\mathbf{q}_{x,d}^t \in \hat{Q}^t(x, d)$, and we may set $\bar{Q}_{x,d}^t = \mathbf{q}_{x,d}^t$. Thus a worst-case (at optimality) distribution $\hat{Q}^ \in \arg \max_{Q \in \text{MAR}} \mathbb{E}_Q[\sum_{t=1}^T C_t^{\hat{\pi}^*}]$ may be given as follows. $\hat{Q}_1^* = \mathbf{q}_{\hat{x}_1^*, \mu}^T$. For $t \in [1, T-1]$ and $\mathbf{d} \in \text{supp}(\hat{Q}_{[t]}^*)$,*

$$\hat{Q}_{t+1|\mathbf{d}}^* = \mathbf{q}_{\hat{x}_{t+1}^*(\mathbf{d}), d_t}^{T-t}.$$

Furthermore, for all $x_0 \leq \beta_\mu^T$, the optimal value of Problem 2.4 equals $\text{Opt}_{\text{MAR}}^T(\mu, U, b)$.

We note that this optimal solution is not necessarily unique, but do not attempt to characterize the family of all possible solutions here. Furthermore, as the optimal policy is essentially given in closed form, it may be computed very efficiently. We now derive a series of results which shed further light on the nature of the optimal policy and worst-case distribution described in Theorem 3. In particular, we prove that if the initial inventory level is sufficiently small, i.e. $x_0 \leq \chi_{\text{MAR}}^T(\mu, U, b)$, then the stochastic inventory and demand process induced by the policy $\hat{\pi}^*$, and sequence of demands distributed as \hat{Q}^* , takes a simpler form which we now describe explicitly. Let $\Lambda^T \triangleq \max(T - \Gamma_\mu^T, 1)$, $D_0^T \triangleq \mu$, $X_0^T \triangleq \frac{\Gamma_\mu^T}{b+T+1} \mu$, and for $t \in [1, \Lambda^T]$,

$$D_t^T \triangleq A_{\min(\Gamma_\mu^T, T-1)}^{T+1-t}, \quad X_t^T \triangleq B_{\Gamma_\mu^T}^{T+1-t}. \quad (2.7)$$

Note that Λ^T represents the first time that D_t^T reaches U .

We now formally relate the corresponding inventory and demand dynamics to an appropriate Markov chain. Consider the following discrete time Markov chain $\{M^T(t) \triangleq (\mathcal{X}_t^T, \mathcal{D}_t^T), t \geq 1\}$, with randomized initial conditions.

$$(\mathcal{X}_1^T, \mathcal{D}_1^T) = \begin{cases} (X_1^T, D_1^T) & \text{w.p. } \frac{\mu}{D_1^T}, \\ (X_1^T, 0) & \text{w.p. } 1 - \frac{\mu}{D_1^T}; \end{cases}$$

for $t \in [2, \Lambda^T - 1]$,

$$(\mathcal{X}_t^T, \mathcal{D}_t^T) = \begin{cases} (\mathcal{X}_{t-1}^T, 0) & \text{w.p. } 1, \text{ if } \mathcal{D}_{t-1}^T = 0, \\ (X_t^T, D_t^T) & \text{w.p. } \frac{D_{t-1}^T}{D_t^T}, \text{ if } \mathcal{D}_{t-1}^T \neq 0, \\ (X_t^T, 0) & \text{w.p. } 1 - \frac{D_{t-1}^T}{D_t^T}, \text{ if } \mathcal{D}_{t-1}^T \neq 0; \end{cases}$$

for $t = \Lambda^T$ (if $\Lambda^T \neq 1$),

$$(\mathcal{X}_t^T, \mathcal{D}_t^T) = \begin{cases} (\mathcal{X}_{t-1}^T, 0) & \text{w.p. 1, if } \mathcal{D}_{t-1}^T = 0, \\ (X_t^T, U) & \text{w.p. } \frac{D_{t-1}^T}{U}, \text{ if } \mathcal{D}_{t-1}^T \neq 0, \\ (X_t^T, 0) & \text{w.p. } 1 - \frac{D_{t-1}^T}{U}, \text{ if } \mathcal{D}_{t-1}^T \neq 0; \end{cases}$$

for $t \geq \min(\Lambda^T + 1, T)$,

$$(\mathcal{X}_t^T, \mathcal{D}_t^T) = \begin{cases} (\mathcal{X}_{\Lambda^T}^T, 0) & \text{w.p. 1, if } \mathcal{D}_{\Lambda^T}^T = 0; \\ (U, U) & \text{w.p. 1, if } \mathcal{D}_{\Lambda^T}^T = U. \end{cases}$$

COROLLARY 1. *If $x_0 \in [0, \chi_{\text{MAR}}^T(\mu, U, b)]$, then one can construct $\hat{\mathbf{D}}^*$ on a common probability space with M^T s.t. $\{(\hat{x}_t^*(\hat{D}_{[t-1]}^*), \hat{D}_t^*), t \in [1, T]\}$ equals $\{(\mathcal{X}_t^T, \mathcal{D}_t^T), t \in [1, T]\}$ w.p.1. Furthermore, $\{D_t^T, t \in [1, \Lambda^T]\}$ and $\{X_t^T, t \in [1, \Lambda^T]\}$ are both monotone increasing, with $X_t^T \leq D_t^T$ for all $t \in [1, \Lambda^T]$ and $D_{\Lambda^T}^T = U$.*

We now give an alternate description of the dynamics described in Corollary 1, explicitly describing the random amount of time until the demand is either 0 or U . Let Z^T denote the r.v., with support on the integers belonging to $[1, \Lambda^T]$, whose corresponding probability measure \mathcal{Z}^T satisfies

$$\mathcal{Z}^T(t) = \begin{cases} (1 - \frac{D_{t-1}^T}{D_t^T}) \frac{\mu}{D_{t-1}^T} & \text{if } t \in [1, \Lambda^T - 1]; \\ \frac{\mu}{D_{t-1}^T} & \text{if } t = \Lambda^T. \end{cases}$$

Also, let Y^T denote a r.v., with support on $\{0, U\}$, independent of Z^T , whose corresponding probability measure \mathcal{Y}^T satisfies

$$\mathcal{Y}^T(x) = \begin{cases} \frac{D_{\Lambda^T-1}^T}{U} & \text{if } x = U; \\ 1 - \frac{D_{\Lambda^T-1}^T}{U} & \text{if } x = 0. \end{cases}$$

COROLLARY 2. *Under the same assumptions, and on the same probability space, as described in Corollary 1, one can also construct Z^T, Y^T s.t. all of the following hold w.p.1. $\{(\mathcal{X}_t^T, \mathcal{D}_t^T), t \in [1, Z^T - 1]\}$ equals $\{(X_t^T, D_t^T), t \in [1, Z^T - 1]\}$. If $Z^T \leq \Lambda^T - 1$, then $(\mathcal{X}_t^T, \mathcal{D}_t^T) = (X_{Z^T}^T, 0)$ for all $t \in [Z^T, T]$. If $Z^T = \Lambda^T$, and $Y^T = U$, then $(\mathcal{X}_{\Lambda^T}^T, \mathcal{D}_{\Lambda^T}^T) = (X_{\Lambda^T}^T, U)$, and $(\mathcal{X}_t^T, \mathcal{D}_t^T) = (U, U)$ for all $t \in [\Lambda^T + 1, T]$. If $Z^T = \Lambda^T$, and $Y^T = 0$, then $(\mathcal{X}_t^T, \mathcal{D}_t^T) = (X_{\Lambda^T}^T, 0)$ for all $t \in [\Lambda^T, T]$. Furthermore, $\hat{D}_t^* = D_t^T I(Z^T > t), t \in [1, \Lambda^T - 1]$; and $\hat{D}_t^* = UI(Z^T = \Lambda^T, Y^T = U), t \in [\Lambda^T, T]$.*

Note that Corollary 2 implies that $\hat{\mathbf{D}}^*$ increases along a sequence of pre-determined constants (i.e. D_t^T) until a random time at which it either drops to 0 or jumps to U , and stays there. Similar martingales have played an important role in studying extremal martingales in other contexts, e.g. certain worst-case distributions which arise in the study of the celebrated prophet inequalities (cf. Hill and Kertz [70]). In light of our results, the appearance of such martingales has an appealing interpretation.

Interpretation 1. *Under our model of demand uncertainty, a robust inventory manager should be most concerned about the possibility that her product becomes obsolete, i.e. demand for her product drops to zero at a random time.*

We again note that the scenario of product obsolescence has been recognized as practically relevant and analyzed extensively in the inventory literature. In contrast to much previous literature on this topic, here our model *did not assume that products would become obsolete* - instead the risks of obsolescence arose naturally from our worst-case analysis. Another interesting feature of our analysis is that even in a worst-case setting, the system may enter a state s.t. for all subsequent time periods, demand becomes perfectly predictable, and can be perfectly matched (at the level U). This is again intuitively appealing, as it coincides with the notion that one's product has reached such a level of popularity that demand predictably attains its maximum. Furthermore, it suggests a certain dichotomy - namely, that either one's product eventually becomes obsolete, or eventually reaches a level of popularity for which demand predictably attains its maximum.

We now present an intuitive description of these dynamics, explained as a game between the “inventory manager” (selector of the policy) and “nature” (selector of the worst-case martingale). Nature reasons that, if her demand in some period t is ever zero, the martingale property ensures that she will order zero in all subsequent periods. This will leave the inventory manager stuck holding all the inventory which she held at the start of period t for the entire remainder of the time horizon, which will incur a large cost. Such a situation is thus desirable from nature's perspective, and one might think that to maximize the probability of this happening, nature will thus maximize the probability that the demand is 0, which would be achieved by putting all probability on 0 and U . However, as Corollaries 1 - 2 show, this is not the case. In fact, the adversary does not put any probability mass at U until time Λ^T , i.e. the last time period in which there are any meaningful dynamics. The reason is that in the martingale-demand model, there is an additional hidden cost for nature associated with putting probability on or near U . In particular, if a demand of U ever occurs, then the martingale property ensures that all future demands must be U as well. However, it is “easy” for an inventory manager to perform well against an adversary that always orders U - in particular, she can simply order up to U in every period, incurring zero cost. In other words, the aforementioned hidden cost to nature comes in the form of a “loss of randomness”, making nature perfectly predictable going forwards. Corollaries 1 - 2 indicate that at optimality, this tradeoff manifests by having nature always put some probability at 0, and some probability on a different quantity D_t^T , which increases monotonically to U as $t \uparrow \Lambda^T$. This “ramping up” can be explained by observing that although higher values for this second point of the support result in a greater “loss of randomness”, the cost of such a loss becomes smaller over time, as the associated window of time during which the adversary is perfectly predictable shrinks accordingly.

Note that Corollary 1 implies that (in a worst-case setting) the inventory manager continues to ramp up production until either the demand drops to zero, or time Λ^T is reached, where the precise manner in which production is ramped up results from an attempt to balance the cost of potential obsolescence (i.e. being stuck holding one's inventory) with the cost of under-ordering. To better understand this ramping up, and more generally the dynamics of an optimal policy against a worst-case distribution for Problem 2.4, we now prove that the Markov chain M^T , properly rescaled, converges weakly (as $T \rightarrow \infty$) to a simple and intuitive limiting process. For an excellent review of the formal definition of weak convergence, and the relevant topological spaces and metrics, we refer the reader to the excellent texts Billingsley [26] and Whitt [128].

Let us define

$$\gamma \triangleq \frac{\mu}{U}, \quad \Lambda^\infty \triangleq 1 - \gamma^{\frac{1}{b}}.$$

Let Z^∞ denote the mixed (i.e. both continuous and discrete components) r.v., with continuous support on $[0, \Lambda^\infty)$, and discrete support on the singleton Λ^∞ , whose corresponding probability measure Z^∞ satisfies

$$Z^\infty(\alpha) = \begin{cases} b(1 - \alpha)^{b-1}d\alpha & \text{if } \alpha \in [0, \Lambda^\infty); \\ \gamma & \text{if } \alpha = \Lambda^\infty. \end{cases}$$

For $\alpha \in [0, \Lambda^\infty]$, where α is again a continuous parameter, let us also define

$$D_\alpha^\infty \triangleq \mu(1-\alpha)^{-b} \quad , \quad X_\alpha^\infty \triangleq \mu\gamma^{\frac{1}{b}}(1-\alpha)^{-(b+1)}.$$

Note that $D_{\Lambda^\infty}^\infty = X_{\Lambda^\infty}^\infty = U$. We now define an appropriate limiting process. Let $\mathcal{M}^\infty(\alpha)_{0 \leq \alpha \leq 1}$ denote the following two-dimensional process, constructed on the same probability space as Z^∞ .

$$\mathcal{M}^\infty(\alpha) = \begin{cases} (X_\alpha^\infty, D_\alpha^\infty) & \text{if } \alpha \in [0, Z^\infty); \\ (X_{Z^\infty}^\infty, 0) & \text{if } \alpha \geq Z^\infty \text{ and } Z^\infty < \Lambda^\infty; \\ (U, U) & \text{if } \alpha \geq Z^\infty \text{ and } Z^\infty = \Lambda^\infty. \end{cases}$$

Then we have the following weak convergence result. For $\alpha \in [0, 1]$, let $\mathcal{M}^T(\alpha) \triangleq M^T(\max(\lfloor \alpha T \rfloor, 1))$.

THEOREM 4. *For all $U, b > 0$ and $\mu \in (0, U)$, the sequence of stochastic processes $\{\mathcal{M}^T(\alpha)_{0 \leq \alpha \leq 1}, T \geq 1\}$ converges weakly to the process $\mathcal{M}^\infty(\alpha)_{0 \leq \alpha \leq 1}$ on the space $D([0, 1], \mathbb{R}^2)$ under the J_1 topology.*

We note that Theorem 4 reveals several qualitative insights into the corresponding dynamics. For example, it provides simple analytical formulas for how (in a worst-case setting) the production and demand both ramp up over time, and demonstrates that the limiting probability that obsolescence ever occurs (again in such a worst-case setting) equals $1 - \gamma$.

A different lens through which one can gain further insight into Theorem 3 follows by observing that for certain values of b and μ , several quantities appearing in our main results simplify considerably.

Observation 1. *When $b = 1$ and $\mu = \frac{kU}{T+1}$ for some $k \in [1, T]$, $A_j^T = \frac{U(j+1)}{T}$ for $j \in [-1, T]$, and $\Lambda^T = T + 1 - k$. For $t \in [1, \Lambda^T]$, $D_t^T = \frac{kU}{T+1-t}$,*

$$X_t^T = \left(\left(1 - \frac{t}{T+2}\right) \left(1 - \frac{t}{T+1}\right) \right)^{-1} \frac{k}{T+2} \frac{k-1}{T+1} U \quad , \quad Opt_{\text{MAR}}^T(\mu, U, b) = \frac{k}{T} \left(1 - \frac{k}{T+1}\right) UT.$$

Furthermore,

$$Z^T(t) = \begin{cases} \frac{1}{T+1} & \text{if } t \in [1, T-k]; \\ \frac{k+1}{T+1} & \text{if } t = T+1-k; \end{cases}$$

and

$$\mathcal{Y}^T(x) = \begin{cases} 1 - \frac{1}{k+1} & \text{if } x = U; \\ \frac{1}{k+1} & \text{if } x = 0. \end{cases}$$

Namely, Z^T is equally likely to take any value in the interval $[1, T-k]$.

Of course, another avenue for more concretely understanding Theorem 3 is to explicitly evaluate all relevant quantities for small values of T , as we now do for the case $T = 2$.

Observation 2.

$$\begin{aligned}
A_{-1}^1 &= 0, & A_0^1 &= U, & A_1^1 &= (b+1)U, & B_{-1}^1 &= 0, & B_0^1 &= 0, & B_1^1 &= U, \\
A_{-1}^2 &= 0, & A_0^2 &= \frac{U}{b+1}, & A_1^2 &= U, & A_2^2 &= (1 + \frac{b}{2})U, & B_{-1}^2 &= 0, & B_0^2 &= 0, \\
B_1^2 &= \frac{U}{b+2}, & B_2^2 &= U, & A_{-1}^3 &= 0, & A_0^3 &= \frac{2U}{(b+1)(b+2)}, & A_1^3 &= \frac{2U}{b+2}, & A_2^3 &= U, \\
A_3^3 &= (1 + \frac{b}{3})U, & B_{-1}^3 &= 0, & B_0^3 &= 0, & B_1^3 &= \frac{2U}{(b+2)(b+3)}, & B_2^3 &= \frac{2U}{b+3}, & B_3^3 &= U.
\end{aligned}$$

$$\left\{ \begin{array}{ll}
\mathbf{q}_{x,\mu}^2(\frac{U}{b+1}) = \frac{U-\mu}{U(1-\frac{1}{b+1})}, & \mathbf{q}_{x,\mu}^2(U) = \frac{\mu - \frac{U}{b+1}}{U(1-\frac{1}{b+1})} & \text{if } \mu \in (\frac{U}{b+1}, U], x \in [0, \frac{U}{b+2}); \\
\mathbf{q}_{x,\mu}^2(0) = 1 - (b+1)\frac{\mu}{U}, & \mathbf{q}_{x,\mu}^2(\frac{U}{b+1}) = (b+1)\frac{\mu}{U} & \text{if } \mu \in (0, \frac{U}{b+1}], x \in [0, \frac{U}{b+2}); \\
\mathbf{q}_{x,\mu}^2(0) = 1 - \frac{\mu}{U}, & \mathbf{q}_{x,\mu}^2(U) = \frac{\mu}{U} & \text{if } \mu \in (0, \frac{U}{b+1}], x \in [\frac{U}{b+2}, U); \\
\mathbf{q}_{x,\mu}^2(0) = 1 - \frac{\mu}{U}, & \mathbf{q}_{x,\mu}^2(U) = \frac{\mu}{U} & \text{if } \mu \in (\frac{U}{b+1}, U], x \in [\frac{U}{b+2}, U); \\
\mathbf{q}_{x,\mu}^2(0) = 1 - \frac{\mu}{U}, & \mathbf{q}_{x,\mu}^2(U) = \frac{\mu}{U} & \text{if } \mu = 0 \text{ or } x = U.
\end{array} \right.$$

$$G_0^2(x, \mu) = (2 - (b+1)(b+2)\frac{\mu}{U})x + 2b\mu, \quad G_1^2(x, \mu) = (2 - (b+2)\frac{\mu}{U})x + b\mu, \quad G_2^2(x, \mu) = 2(1 - \frac{\mu}{U})x.$$

$$\Gamma_\mu^2 = \begin{cases} 0 & \text{if } \mu \in [0, \frac{2U}{(b+1)(b+2)}], \\ 1 & \text{if } \mu \in (\frac{2U}{(b+1)(b+2)}, \frac{2U}{b+2}], \\ 2 & \text{if } \mu \in (\frac{2U}{b+2}, U]. \end{cases}, \quad \Lambda^2 = \begin{cases} 2 & \text{if } \mu \in [0, \frac{2U}{(b+1)(b+2)}], \\ 1 & \text{if } \mu \in (\frac{2U}{(b+1)(b+2)}, U]. \end{cases}$$

$$\chi_{\text{MAR}}^2 = \begin{cases} 0 & \text{if } \mu \in [0, \frac{2U}{(b+1)(b+2)}], \\ \frac{U}{b+2} & \text{if } \mu \in (\frac{2U}{(b+1)(b+2)}, \frac{2U}{b+2}], \\ U & \text{if } \mu \in (\frac{2U}{b+2}, U]. \end{cases}, \quad \text{Opt}_{\text{MAR}}^2 = \begin{cases} 2b\mu & \text{if } \mu \in [0, \frac{2U}{(b+1)(b+2)}], \\ \frac{2U}{b+2} + (b-1)\mu & \text{if } \mu \in (\frac{2U}{(b+1)(b+2)}, \frac{2U}{b+2}], \\ 2(U - \mu) & \text{if } \mu \in (\frac{2U}{b+2}, U]. \end{cases}$$

$$D_1^2 = \begin{cases} \frac{U}{b+1} & \text{if } \mu \in [0, \frac{2U}{(b+1)(b+2)}], \\ U & \text{if } \mu \in (\frac{2U}{(b+1)(b+2)}, U]. \end{cases}, \quad D_2^2 = \begin{cases} U & \text{if } \mu \in [0, \frac{2U}{(b+1)(b+2)}]. \end{cases}$$

$$X_1^2 = \begin{cases} 0 & \text{if } \mu \in [0, \frac{2U}{(b+1)(b+2)}], \\ \frac{U}{b+2} & \text{if } \mu \in (\frac{2U}{(b+1)(b+2)}, \frac{2U}{b+2}], \\ U & \text{if } \mu \in (\frac{2U}{b+2}, U]. \end{cases}, \quad X_2^2 = \begin{cases} 0 & \text{if } \mu \in [0, \frac{2U}{(b+1)(b+2)}]. \end{cases}$$

$$\begin{cases} \mathcal{Z}^2(1) = 1 - (b+1)\frac{\mu}{U}, & \mathcal{Z}^2(2) = (b+1)\frac{\mu}{U} & \text{if } \mu \in [0, \frac{2U}{(b+1)(b+2)}]; \\ \mathcal{Z}^2(1) = 1 & & \text{if } \mu \in (\frac{2U}{(b+1)(b+2)}, U]. \end{cases}$$

$$\begin{cases} \mathcal{Y}^2(U) = \frac{1}{b+1}, & \mathcal{Y}^2(0) = 1 - \frac{1}{b+1} & \text{if } \mu \in [0, \frac{2U}{(b+1)(b+2)}]; \\ \mathcal{Y}^2(U) = \frac{\mu}{U}, & \mathcal{Y}^2(0) = 1 - \frac{\mu}{U} & \text{if } \mu \in (\frac{2U}{(b+1)(b+2)}, U]. \end{cases}$$

We note that Observation 2 brings another interesting feature of Theorem 3 to light, namely the fact that $\chi_{\text{MAR}}^T(\mu, U, b)$ is not monotone increasing in b . This is surprising, since one would expect that as the backlogging penalty increases, one would wish to stock higher inventory levels. We also note that for fixed T, μ, U , $\chi_{\text{IND}}^T(\mu, U, b)$ is monotone increasing in b , i.e. such a non-monotonicity does not manifest in the independent-demand model. Furthermore, Theorem 4, and the monotonicity (in b) of $X_0^\infty = \mu\gamma^{\frac{1}{b}}$, implies that this non-monotonicity does not manifest “in the limit”.

We conclude by presenting several comparative results between the independent-demand and martingale-demand models. First, we prove that the expected cost incurred by a minimax optimal policy is highest under the independent-demand model, among all possible dependency structures, whenever $x_0 = 0$. In particular, in this case the expected cost incurred under the independent-demand model is *at least* the expected cost incurred under the martingale-demand

model. Let \mathbf{GEN}_T denote the family of all T -dimensional measures Q s.t. $Q_t \in \mathfrak{M}(\mu)$ for all t , and $\text{Opt}_{\mathbf{GEN}}^T(\mu, U, b) \triangleq \inf_{\pi \in \Pi} \sup_{Q \in \mathbf{GEN}_T} \mathbb{E}_Q[\sum_{t=1}^T C_t^\pi]$ under the initial condition $x_0 = 0$. Then we prove the following comparative result, and defer the proof to the Technical Appendix in Section 6.

THEOREM 5. *For all strictly positive $U, b \in \mathcal{Q}^+, T \geq 1$, and $\mu \in [0, U]_{\mathcal{Q}}$, $\text{Opt}_{\mathbf{GEN}}^T(\mu, U, b) = \text{Opt}_{\mathbf{IND}}^T(\mu, U, b)$.*

Although Theorem 5 implies that $\frac{\text{Opt}_{\mathbf{MAR}}^T(\mu, U, b)}{\text{Opt}_{\mathbf{IND}}^T(\mu, U, b)} \leq 1$, it is interesting to further understand the behavior of this ratio, in the spirit of the so-called price of correlations introduced in Agrawal et al. [2]. Indeed, we prove the following, which provides a simple analytical form for this ratio as $T \rightarrow \infty$, again deferring the proof to the Technical Appendix in Section 6.

THEOREM 6. *For all strictly positive $U, b \in \mathcal{Q}^+$, and $\mu \in (0, U)_{\mathcal{Q}}$,*

$$\lim_{T \rightarrow \infty} \frac{\text{Opt}_{\mathbf{MAR}}^T(\mu, U, b)}{\text{Opt}_{\mathbf{IND}}^T(\mu, U, b)} = \begin{cases} 1 - \gamma^{\frac{1}{b}} & \text{if } \mu \leq \frac{U}{b+1}, \\ \frac{(1-\gamma^{\frac{1}{b}})b\mu}{U-\mu} & \text{if } \mu > \frac{U}{b+1}. \end{cases}$$

We also note the following corollary, which follows immediately from Theorem 6.

COROLLARY 3. *Suppose $U \in \mathcal{Q}^+, \mu = \frac{U}{2}$, and $b = 1$. Then*

$$\lim_{T \rightarrow \infty} \frac{\text{Opt}_{\mathbf{MAR}}^T(\mu, U, b)}{\text{Opt}_{\mathbf{IND}}^T(\mu, U, b)} = \frac{1}{2}.$$

One setting which arises in various inventory applications is that in which $b \gg h$, which motivates considering what our results say about the regime in which $b \rightarrow \infty$. In particular, the following may be easily verified by combining our main result Theorem 3 with a straightforward calculation.

COROLLARY 4. *For all strictly positive $U \in \mathcal{Q}^+, \mu \in [0, U]_{\mathcal{Q}}, T \geq 1$, and $b \in \mathcal{Q}^+$ s.t. $b > (\frac{U}{\mu} - 1)T$, it holds that $\chi_{\mathbf{MAR}}^T(\mu, U, b) = \chi_{\mathbf{IND}}(\mu, U, b) = U$, and $\frac{\text{Opt}_{\mathbf{MAR}}^T(\mu, U, b)}{\text{Opt}_{\mathbf{IND}}^T(\mu, U, b)} = 1$.*

Intuitively, Corollary 4 indicates that with all other parameters fixed, if b is sufficiently large, then in both scenarios the inventory manager orders up to U in the first period out of fear of incurring large backlogging costs. Furthermore, although the demand dynamics are different in the two models in this case, they both result in the same minimax expected cost. Interestingly, when one considers the case of large b in the asymptotic setting of Theorem 6 (in which T is also very large), a different behavior emerges. Namely, by letting $b \rightarrow \infty$ in Theorem 6, one may conclude the following after a straightforward calculation.

COROLLARY 5. *For all strictly positive $U \in \mathcal{Q}^+$, and $\mu \in (0, U)_{\mathcal{Q}}$,*

$$\lim_{b \rightarrow \infty} \lim_{T \rightarrow \infty} \frac{\text{Opt}_{\mathbf{MAR}}^T(\mu, U, b)}{\text{Opt}_{\mathbf{IND}}^T(\mu, U, b)} = \frac{\log(\gamma^{-1})}{\gamma^{-1} - 1}.$$

Combining Corollaries 4 and 5, we conclude that in this case an interchange of limits does not hold. Namely, as Corollary 4 implies that $\lim_{T \rightarrow \infty} \lim_{b \rightarrow \infty} \frac{\text{Opt}_{\mathbf{MAR}}^T(\mu, U, b)}{\text{Opt}_{\mathbf{IND}}^T(\mu, U, b)} = 1$, here the order in which one lets b, T get large makes a fundamental difference. Indeed, as it is easily verified that $\lim_{b \rightarrow \infty} (b\Lambda^\infty) = \log(\gamma^{-1})$, while $D_0^\infty = \mu$ irregardless of b , what happens in this second setting as b gets large is that D_α^∞ ramps up very quickly from μ to U , where the time required to ramp up behaves (asymptotically) as $\frac{\log(\gamma^{-1})}{b}T$ (remember that as large as b may be, T is always much larger due

to the order of limits here). However, over this relatively short time horizon (short as a fraction of T , still large in an absolute sense as T is very large), a non-trivial backlogging penalty may be incurred since b is so large. It can be shown that these effects result in a non-trivial limiting dynamics under a proper rescaling, although we do not present a complete analysis here for the sake of brevity.

As a final observation, we note that the comparative results of Theorem 5 are sensitive to the initial conditions, and in fact for different values of x_0 the situation may be reversed. For example, we have the following result, whose proof we similarly defer to the Technical Appendix in Section 6.

Observation 3. *For all strictly positive $U, b \in \mathcal{Q}^+, T \geq 1$, and $\mu \in [0, U]_{\mathcal{Q}}$, if $x_0 = U$ and $\frac{\mu}{U} < \frac{1}{b+1}$, then $\lim_{T \rightarrow \infty} \frac{\inf_{\pi \in \Pi} \sup_{Q \in \text{MAR}} \mathbb{E}_Q[\sum_{t=1}^T C_t^{\pi}]}{\inf_{\pi \in \Pi} \sup_{Q \in \text{IND}} \mathbb{E}_Q[\sum_{t=1}^T C_t^{\pi}]} = \frac{U-\mu}{b\mu} > 1$.*

3. Proof of Theorem 3 and Corollaries 1 - 2

In this section, we complete the proof of our main result, Theorem 3, which yields an explicit solution to Problem 2.4, as well as Lemma 2, and Corollaries 1 - 2. We proceed by induction, and note that our analysis combines ideas from convex analysis with the theory of martingales. We begin by making several additional definitions, and proving several auxiliary results.

For $x, \mu \in [0, U]$, and $j \in [-1, T-1]$, let

$$F_j^T(x, \mu) \triangleq -bx + (b+T)B_{j+1}^T + (Tb - (b+1)(j+1))\mu.$$

Also, let us define

$$\mathfrak{g}^T(x, \mu) \triangleq \begin{cases} F_j^T(x, \mu) & \text{if } \mu \in (A_j^T, A_{j+1}^T], x \in [0, B_{j+1}^T); \\ G_k^T(x, \mu) & \text{if } \mu \in (A_j^T, A_{j+1}^T], x \in [B_k^T, B_{k+1}^T), k \geq j+1; \\ G_k^T(x, 0) & \text{if } \mu = 0, x \in [B_k^T, B_{k+1}^T); \\ G_T^T(U, \mu) & \text{if } x = U. \end{cases}$$

For later proofs, it will be convenient to note that \mathfrak{g} may be equivalently expressed as follows. Let

$$\Upsilon_x^T \triangleq \begin{cases} T & \text{if } x = U, \\ j & \text{if } x \in [B_j^{T+1}, B_{j+1}^{T+1}); \end{cases}$$

and note that $x \in [B_{\Upsilon_x^T}^T, B_{\Upsilon_x^T+1}^T)$ for all $T \geq 2$ and $x \in [0, U)$, while $U = B_{\Upsilon_U^T}^T$. Noting that $G_T^T(U, \mu) = G_{T-1}^T(U, \mu)$ implies

$$\mathfrak{g}^T(x, \mu) = \begin{cases} F_{\Upsilon_x^T-1}^T(x, \mu) & \text{if } x < B_{\Upsilon_x^T}^T; \\ G_{\Upsilon_x^T}^T(x, \mu) & \text{if } x \geq B_{\Upsilon_x^T}^T. \end{cases}$$

We now prove several properties of \mathfrak{g}^T . It will first be useful to review a well-known sufficient condition for convexity of non-differentiable functions. Recall that for a one-dimensional function $f(x)$, the right-derivative of f evaluated at x_0 , which we denote by $\partial_x^+ f(x_0)$, equals $\lim_{h \downarrow 0} \frac{f(x_0+h) - f(x_0)}{h}$. When this limit exists, we say that f is right-differentiable at x_0 . Then the following sufficient condition for convexity is stated in Royden and Fitzpatrick [107] Section 5, Proposition 18.

LEMMA 3 (Royden and Fitzpatrick [107]). *A one-dimensional function $f(x)$, which is continuous and right-differentiable on an open interval (a, b) with non-decreasing right-derivative on (a, b) , is convex on (a, b) .*

Then we may derive the following properties of \mathbf{g}^T and other relevant quantities, whose proofs we defer to the Technical Appendix in Section 6.

LEMMA 4. $\mathbf{g}^T(x, d)$ is a continuous and convex function of x on $(0, U)$, and a right (left) continuous function of x at 0 (U).

LEMMA 5. $A_j^T < A_{j+1}^{T+1} < A_{j+1}^T$ for $j \in [0, T-2]$. It follows that for all $T \geq 2$ and $d \in [0, U]$, $\Gamma_d^T \in \{\Gamma_d^{T-1}, \Gamma_d^{T-1} + 1\}$.

LEMMA 6. $\beta_d^T \in \arg \min_{z \in [0, U]} \mathbf{g}^T(z, d)$.

To simplify various notations and concepts before completing the proof of Theorem 3, it will be useful to make several additional definitions, and prove a few more relevant preliminary results. For $x, d \in [0, U]$, and $j \in [0, T]$, let

$$\overline{F}_j^T(x, d) \triangleq Tx + ((b-1)T - bj - \frac{b+T}{A_j^T}x)d + \frac{b+T}{A_j^T}d^2;$$

and for $j \in [-1, T-1]$, let

$$\overline{G}_j^T(d) \triangleq TB_{j+1}^T + (Tb - (b+1)(j+1))d.$$

Note that β_d^T is monotone increasing in d , with $\beta_0^T = 0$, and $\beta_U^T = U$. It follows that for all $x \in [0, U]$,

$$z_x^T \triangleq \inf\{d \geq 0 \quad \text{s.t.} \quad \beta_d^T \geq x - d\},$$

is well defined,

$$z_x^T \leq x \quad \text{for} \quad x \in [0, U], \tag{3.8}$$

and $z_x^T = 0$ iff $x = 0$. For $x, d \in [0, U]$, let us define

$$\overline{\mathbf{g}}^T(x, d) \triangleq \begin{cases} \overline{F}_j^T(x, d) & \text{if } d \in [0, z_x^T] \cap (x - B_{j+1}^T, x - B_j^T], \quad j \in [0, T-1]; \\ \overline{G}_j^T(d) & \text{if } d \in [z_x^T, U] \cap (A_j^{T+1}, A_{j+1}^{T+1}], \quad j \in [-1, T-1]; \\ \overline{G}_{-1}^T(0) & \text{if } d = 0, x = 0; \\ \overline{F}_{T-1}^T(U, 0) & \text{if } d = 0, x = U. \end{cases}$$

For later proofs, it will be convenient to note that $\overline{\mathbf{g}}$ may be equivalently expressed as follows.

$$\overline{\mathbf{g}}^T(x, d) = \begin{cases} \overline{F}_{\Gamma_{x-d}^T}^T(x, d) & \text{if } d < z_x^T; \\ \overline{G}_{\Gamma_d^T}^T(d) & \text{if } d \geq z_x^T. \end{cases}$$

In that case, we may derive the following properties of $\overline{\mathbf{g}}^T(x, d)$, whose proofs we defer to the Technical Appendix in Section 6.

LEMMA 7. For all $T \geq 1$, $x, d \in [0, U]$, $\overline{\mathbf{g}}^T(x, d) = \mathbf{g}^T\left(\max(\beta_d^T, x - d), d\right)$.

LEMMA 8. For each fixed $x \in [0, U]$, $\overline{\mathbf{g}}^T(x, d)$ is a continuous function of d on $(0, U)$, and a right (left) - continuous function of d at 0 (U). Also, $\overline{\mathbf{g}}^T(x, d)$ is a convex function of d on $(0, z_x^T)$, and a concave function of d on (z_x^T, U) .

Also, for $x, d \in [0, U]$, let us define $\eta_d \triangleq x - d$, $\alpha_x^T \triangleq A_{\Gamma_x^T}^T$, $\zeta_x^T \triangleq \Gamma_{z_x^T}^T$, $\mathcal{A}_x^T \triangleq A_{\zeta_x^T}^T$, $\mathfrak{N}_x^T \triangleq A_{\Upsilon_x^T}^T$,

$$\mathfrak{f}^T(x, d) \triangleq b(d - x)_+ + (x - d)_+ + \bar{\mathfrak{g}}^{T-1}(x, d).$$

Note that for all $j \in [-1, T - 1]$,

$$\begin{aligned} B_{j+1}^T &= \frac{j+1}{b+T} A_{j+1}^T \\ &\leq \frac{j+1}{b+j+1} A_{j+1}^T = A_j^T. \end{aligned} \quad (3.9)$$

Combining with (3.8) and a straightforward contradiction argument, we conclude that for all $x \in [0, U]$ and $T \geq 2$,

$$\Upsilon_x^T \geq \Gamma_x^T \geq \zeta_x^T \quad , \quad \text{and} \quad \mathfrak{N}_x^T \geq \alpha_x^T \geq \mathcal{A}_x^T. \quad (3.10)$$

In that case, we may derive the following structural properties of \mathfrak{f}^T , whose proofs we defer to the Technical Appendix in Section 6.

LEMMA 9. *For each fixed $x \in [0, U]$, $\mathfrak{f}^T(x, d)$ is a continuous function of d on $(0, U)$, and a right (left) - continuous function of d at 0 (U). Also, for all $j \in [-1, T - 2]$, $\mathfrak{f}^T(x, d)$ is a convex function of d on (A_j^T, A_{j+1}^T) . Furthermore, $\mathfrak{f}^T(x, d)$ is a convex function of d on $(0, \mathcal{A}_x^T)$, and a concave function of d on (α_x^T, U) .*

We now introduce one last preliminary result, which will provide a mechanism for certifying a distribution as the solution to a given distributionally robust optimization problem, whose proof we again defer to the Technical Appendix in Section 6.

LEMMA 10. *Suppose $f: [0, U] \rightarrow \mathbb{R}$ is any bounded real-valued function with domain $[0, U]$. Suppose $0 \leq L \leq \mu < R \leq U$, and the linear function $\eta(d) \triangleq \frac{f(R)-f(L)}{R-L}d + \frac{Rf(L)-Lf(R)}{R-L}$, i.e. the line intersecting f at the points L and R , satisfies $\eta(d) \geq f(d)$ for all $d \in [0, U]$, i.e. lies above f on $[0, U]$. Then the measure q s.t. $q(L) = \frac{R-\mu}{R-L}$, $q(R) = \frac{\mu-L}{R-L}$, belongs to $\arg \max_{Q \in \mathfrak{M}(\mu)} \mathbb{E}_Q[f(D)]$.*

With all of the above notations, definitions, and preliminary results in place, we now proceed with the proof of Theorem 3, which will follow from the following result.

THEOREM 7. *For all $x, d \in [0, U]_{\mathcal{Q}}$ and $T \geq 1$, $\mathfrak{g}^T(x, d) = \hat{\mathfrak{g}}^T(x, d)$, and $\mathfrak{q}_{x,d}^T \in \hat{\mathcal{Q}}^T(x, d)$.*

Proof: We proceed by induction, beginning with the base case $T = 1$. It follows from Observation 2 that for all $x, \mu \in [0, U]_{\mathcal{Q}}$, $\Gamma_{\mu}^0 = \Upsilon_x^0 = B_{\Gamma_{\mu}^0}^1 = 0$, $\Gamma_{\mu}^1 = I(\mu > \frac{U}{b+1})$, $\beta_{\mu}^1 = U \times I(\mu > \frac{U}{b+1})$, $\mathfrak{g}^1(x, \mu) = G_0^1(x, \mu) = (1 - \frac{b+1}{U}\mu)x + b\mu$, and $\mathfrak{q}_{x,\mu}^1(0) = 1 - \frac{\mu}{U}$, $\mathfrak{q}_{x,\mu}^1(U) = \frac{\mu}{U}$. The desired result then follows directly from the results of Shapiro [115], and we omit the details.

Now, suppose that the result is true for all $s \in [1, T - 1]$ for some $T \geq 2$. We now prove that the results hold also for T . It follows from the induction hypothesis, and Lemmas 4, 6, and 7 that for all $x, d \in [0, U]_{\mathcal{Q}}$,

$$\hat{V}^{T-1}(x, d) = \mathfrak{g}^{T-1}\left(\max(\beta_d^{T-1}, x), d\right) \quad , \quad \hat{V}^{T-1}(x-d, d) = \bar{\mathfrak{g}}^{T-1}(x, d) \quad , \quad \text{and} \quad \hat{f}^T(x, d) = \mathfrak{f}^T(x, d). \quad (3.11)$$

From definitions, to prove the desired induction, it thus suffices to demonstrate that

$$\mathfrak{g}^T(x, d) = \sup_{Q \in \mathfrak{M}(d)} \mathbb{E}_Q[\mathfrak{f}^T(x, D)] \quad , \quad \text{and} \quad \mathfrak{q}_{x,d}^T \in \arg \max_{Q \in \mathfrak{M}(d)} \mathbb{E}_Q[\mathfrak{f}^T(x, D)]. \quad (3.12)$$

We now use Lemma 9 to explicitly construct (for each fixed value of x) a family \mathcal{F} of lines $\{L_i\}$, s.t. each line L_i lies above $\mathfrak{f}^T(x, d)$ for all $d \in [0, U]$, each line L_i intersects $\mathfrak{f}^T(x, d)$ at exactly two points p_i^1, p_i^2 , and for each $\mu \in [0, U]$ there exists $L_i \in \mathcal{F}$ s.t. $\mu \in [p_i^1, p_i^2]$. Our construction will ultimately allow us to apply Lemma 10, explicitly solve the distributionally robust optimization problem $\sup_{Q \in \mathfrak{M}(d)} \mathbb{E}_Q[\mathfrak{f}^T(x, D)]$, and complete the proof. We begin by explicitly constructing the family of lines \mathcal{F} . For $d \in \mathbb{R}$, let us define

$$\mathcal{K}^T(x, d) \triangleq \frac{\mathfrak{f}^T(x, \aleph_x^T) - Tx}{\aleph_x^T} d + Tx;$$

and for $j \in [-1, T-2]$,

$$\mathcal{L}_j^T(x, d) \triangleq \frac{\mathfrak{f}^T(x, A_{j+1}^T) - \mathfrak{f}^T(x, A_j^T)}{A_{j+1}^T - A_j^T} d + \frac{A_{j+1}^T \mathfrak{f}^T(x, A_j^T) - A_j^T \mathfrak{f}^T(x, A_{j+1}^T)}{A_{j+1}^T - A_j^T}.$$

Noting that $\overline{F}_j^{T-1}(x, 0) = (T-1)x$ for all j , it follows from Lemma 7 that

$$\mathfrak{f}^T(x, 0) = Tx. \quad (3.13)$$

It may be easily verified, using (3.13), that \mathcal{K}^T defines the unique line passing through the (x, y) coordinates $(0, \mathfrak{f}^T(x, 0))$ and $(\aleph_x^T, \mathfrak{f}^T(x, \aleph_x^T))$; and \mathcal{L}_j^T defines the unique line passing through the (x, y) co-ordinates $(A_j^T, \mathfrak{f}^T(x, A_j^T))$ and $(A_{j+1}^T, \mathfrak{f}^T(x, A_{j+1}^T))$. Then we have the following result, showing that \mathcal{K}^T lies above \mathfrak{f}^T , and that \mathcal{L}_ℓ^T lies above \mathfrak{f}^T for certain values of ℓ , whose proof we defer to the Technical Appendix in Section 6.

LEMMA 11. *For each fixed $x \in [0, U]$, $\mathcal{K}^T(x, d) \geq \mathfrak{f}^T(x, d)$ for all $d \in [0, U]$. Also, for all $\ell \in [\Upsilon_x^{T-1}, T-2]$, $\mathcal{L}_\ell^T(x, d) \geq \mathfrak{f}^T(x, d)$ for all $d \in [0, U]$.*

That $\mathfrak{q}_{x,d}^T \in \arg \max_{Q \in \mathfrak{M}(d)} \mathbb{E}_Q[\mathfrak{f}^T(x, D)]$ follows from Lemmas 10 and 11, definitions, and a straightforward case analysis, the details of which we omit. We now prove that $\mathfrak{g}^T(x, d) = \sup_{Q \in \mathfrak{M}(d)} \mathbb{E}_Q[\mathfrak{f}^T(x, D)]$, and proceed by a case analysis. Let $j = \Gamma_d^{T-1}$ and $k = \Upsilon_x^{T-1}$. First, suppose $d \in [0, A_k^T]$. In light of Lemmas 10 and 11, in this case it suffices to demonstrate that

$$\mathfrak{g}^T(x, d) = \mathcal{K}^T(x, d). \quad (3.14)$$

In this case, the left-hand side of (3.14) equals

$$G_k^T(x, d) = (T - \frac{b+T}{A_k^T} d)x + (T-k)bd. \quad (3.15)$$

Alternatively, from (6.59), the right-hand side of (3.14) equals

$$(b(T-k) - \frac{(b+T)x}{A_k^T})d + Tx. \quad (3.16)$$

Noting that (3.15) equals (3.16) completes the proof in this case.

Alternatively, suppose $d > A_k^T$. In this case it suffices to demonstrate that

$$\mathfrak{g}^T(x, d) = \mathcal{L}_{j-1}^T(x, d). \quad (3.17)$$

Note that the left-hand side of (3.17) equals

$$F_{j-1}^T(x, d) = -bx + (b+T)B_j^T + (Tb - (b+1)j)d. \quad (3.18)$$

Alternatively, it follows from (3.8) and (3.9) that the right-hand side of (3.17) equals

$$\begin{aligned} & \frac{\mathfrak{f}^T(x, A_j^T) - \mathfrak{f}^T(x, A_{j-1}^T)}{A_j^T - A_{j-1}^T} d + \frac{A_j^T \mathfrak{f}^T(x, A_{j-1}^T) - A_{j-1}^T \mathfrak{f}^T(x, A_j^T)}{A_j^T - A_{j-1}^T} \\ &= b(d-x) + \frac{\overline{G}_{j-1}^{T-1}(A_j^T) - \overline{G}_{j-2}^{T-1}(A_{j-1}^T)}{A_j^T - A_{j-1}^T} d + \frac{A_j^T \overline{G}_{j-2}^{T-1}(A_{j-1}^T) - A_{j-1}^T \overline{G}_{j-1}^{T-1}(A_j^T)}{A_j^T - A_{j-1}^T}. \end{aligned} \quad (3.19)$$

We now simplify (3.19). Note that $\overline{G}_{j-1}^{T-1}(A_j^T) - \overline{G}_{j-2}^{T-1}(A_{j-1}^T)$ equals

$$\begin{aligned} & (T-1)(B_j^{T-1} - B_{j-1}^{T-1}) + \left((T-1)b - (b+1)j \right) (A_j^T - A_{j-1}^T) - (b+1)A_{j-1}^T \\ &= \left((T-1)b - (b+1)j \right) (A_j^T - A_{j-1}^T) + (T-1) \left(\frac{j}{T-1} A_j^T - \frac{j-1}{T-1} A_{j-1}^T \right) - (b+1)A_{j-1}^T \\ &= \left((T-1)b - (b+1)j \right) (A_j^T - A_{j-1}^T), \end{aligned} \quad (3.20)$$

and $A_j^T \overline{G}_{j-2}^{T-1}(A_{j-1}^T) - A_{j-1}^T \overline{G}_{j-1}^{T-1}(A_j^T)$ equals

$$\begin{aligned} & (T-1) (A_j^T B_{j-1}^{T-1} - A_{j-1}^T B_j^{T-1}) + (b+1)A_{j-1}^T A_j^T \\ &= (T-1) \left(A_j^T A_{j-1}^T \frac{j-1}{T-1} - A_{j-1}^T A_j^T \frac{j}{T-1} \right) + (b+1)A_{j-1}^T A_j^T \\ &= bA_{j-1}^T A_j^T = jA_j^T (A_j^T - A_{j-1}^T). \end{aligned} \quad (3.21)$$

Plugging (3.20) and (3.21) into (3.19), and comparing to (3.18), completes the proof of (3.12), and the desired induction. \square

With Theorem 7 in hand, we now complete the proof of our main result Theorem 3, as well as Corollaries 1 and 2.

Proof: [Proof of Theorem 3] Lemma 2 follows from Theorem 7 and Lemmas 4 and 6, and it follows from definitions and a straightforward case analysis (the details of which we omit) that $\text{Opt}_{\text{MAR}}^T(\mu, U, b) = \mathfrak{g}^T(\beta_\mu^T, \mu)$. Theorem 3 then follows by combining Lemmas 4 and 6 with Theorem 7, (3.11), Lemma 2, and Theorem 2. \square

Proof: [Proof of Corollaries 1 and 2] First, note that if $\Gamma_\mu^T = T$, then $\chi_{\text{MAR}}^T(\mu, U, b) = B_T^T = A_{T-1}^T = X_1^T = D_1^T = U$, proving the desired result in this case. Next, suppose $\Gamma_\mu^T \leq T-1$. In light of the definitions of X_t^T and D_t^T , i.e. (2.7), the already established monotonicities of $\{A_j^T, j \in [-1, T-1]\}$ and $\{B_j^T, j \in [-1, T]\}$, and the martingale property, it suffices to establish that for all $t \in [1, \Lambda^T]$: 1. $\chi_{\text{MAR}}^{T+1-t}(D_{t-1}^T, U, b) = X_t^T$, 2. $\exists j \in [-1, T-1], k \geq j+1$ s.t. $D_{t-1}^T \in (A_j^{T+1-t}, A_{j+1}^{T+1-t}]$, $X_t^T \in [B_k^{T+1-t}, B_{k+1}^{T+1-t})$, $A_k^{T+1-t} = D_t^T$.

First, suppose $t=1$. In this case, $D_{t-1}^T = \mu \in (A_{\Gamma_\mu^{T-1}-1}^T, A_{\Gamma_\mu^{T-1}}^T]$,

$$\chi_{\text{MAR}}^{T+1-t}(D_{t-1}^T, U, b) = \chi_{\text{MAR}}^T(\mu, U, b) = B_{\Gamma_\mu^T}^T = X_t^T,$$

and $D_t^T = A_{\Gamma_\mu^T}^T$. Note that since A_j^τ is monotone decreasing in τ , it follows that Γ_μ^τ is monotone increasing in τ , and thus $\Gamma_\mu^T \geq \Gamma_\mu^{T-1}$. Thus we find that 1. holds, and 2. holds with $j = \Gamma_\mu^{T-1} - 1, k = \Gamma_\mu^T$.

Next, suppose $t \in [2, \Lambda^T]$. In this case, $D_{t-1}^T = z \triangleq A_{\Gamma_\mu^{T+2-t}}^T$. Note that

$$z \in (A_{\Gamma_\mu^{T-1}}^{T+2-t}, A_{\Gamma_\mu^T}^{T+2-t}), \quad (3.22)$$

and thus $\Gamma_z^{T+1-t} = \Gamma_\mu^T$. Furthermore,

$$\chi_{\text{MAR}}^{T+1-t}(D_{t-1}^T, U, b) = B_{\Gamma_z^{T+1-t}}^{T+1-t} = B_{\Gamma_\mu^T}^{T+1-t} = X_t^T.$$

In addition, the fact that A_j^τ is monotone decreasing in τ and monotone increasing in j , combined with (3.22), implies that $z \in (A_{j'-1}^{T+1-t}, A_{j'}^{T+1-t}]$ for some $j' \leq \Gamma_\mu^T$. Thus we find that 1. holds, and 2. holds with $j = j' - 1, k = \Gamma_\mu^T$, completing the proof. \square

4. Asymptotic analysis and proof of Theorem 4

In this section, we complete the proof of Theorem 4, and begin with a lemma which asserts the (uniform) convergence of several one-dimensional functions to appropriate limits. As all of these convergences follow from repeated application of elementary bounds such as Taylor's Inequality for the exponential function, and the continuous differentiability of all relevant functions over appropriate domains, we omit the details. For $\alpha \in [0, 1)$, let $\mathbf{f}(\alpha) \triangleq b(1-\alpha)^{b-1}$, where we note that $\mathbf{f}(\alpha) = Z^\infty(\alpha)$ for $\alpha \in [0, \Lambda^\infty)$.

LEMMA 12. *For all $\mu \in (0, U)$,*

$$\lim_{T \rightarrow \infty} (T^{-1} \Gamma_\mu^T) = \gamma^{\frac{1}{b}}, \quad \lim_{T \rightarrow \infty} A_{\Gamma_\mu^T}^{T+1} = \mu, \quad \lim_{T \rightarrow \infty} (T^{-1} \Lambda^T) = \Lambda^\infty;$$

$$\limsup_{T \rightarrow \infty} \sup_{t \in [1, \Lambda^T]} |X_t^T - X_{\frac{t}{T}}^\infty| = \limsup_{T \rightarrow \infty} \sup_{t \in [1, \Lambda^T]} |D_t^T - D_{\frac{t}{T}}^\infty| = \limsup_{T \rightarrow \infty} \sup_{t \in [1, \Lambda^T-1]} |T \mathcal{Z}^T(t) - \mathbf{f}(\frac{t}{T})| = 0.$$

For all $T \geq 1, \alpha \geq 0$, let $m_\alpha(T) \triangleq \min(\lfloor \alpha T \rfloor, \Lambda^T - 1)$, and $m_\alpha \triangleq \min(\alpha, \Lambda^\infty)$. For $\alpha_1, \alpha_2 \geq 0$ s.t. $\alpha_1 < m_{\alpha_2}$, let us define $X_{\alpha_1, \alpha_2}^T \triangleq X_{Z^T}^T I([\alpha_1 T] + 1 \leq Z^T \leq m_{\alpha_2}(T))$, and $X_{\alpha_1, \alpha_2}^\infty \triangleq X_{Z^\infty}^\infty I(\alpha_1 \leq Z^\infty < m_{\alpha_2})$. Also, for $\mu \in (0, U)$, note that \mathbf{f} is continuous on $[0, \Lambda^\infty]$, and let $\bar{\mathbf{f}} \triangleq \sup_{\alpha \in [0, \Lambda^\infty]} \mathbf{f}(\alpha) < \infty$. In light of Lemma 12, it follows that

$$\limsup_{T \rightarrow \infty} \sup_{t \in [1, \Lambda^T-1]} (T \mathcal{Z}^T(t)) = \bar{\mathbf{f}}. \quad (4.23)$$

Let us also recall that X^∞ and D^∞ are continuous and strictly increasing on $[0, \Lambda^\infty]$.

Before embarking on the proof of Theorem 4, it will be useful to first establish the weak convergence of X_{α_1, α_2}^T to $X_{\alpha_1, \alpha_2}^\infty$.

LEMMA 13. *For $\mu \in (0, U), \alpha_1, \alpha_2 \geq 0$ s.t. $\alpha_1 < m_{\alpha_2}$, X_{α_1, α_2}^T converges weakly to $X_{\alpha_1, \alpha_2}^\infty$. Furthermore, for any bounded uniformly continuous function $H: [0, U] \rightarrow \mathbb{R}$,*

$$\lim_{T \rightarrow \infty} \sum_{k=\lfloor \alpha_1 T \rfloor + 1}^{m_{\alpha_2}(T)} H(X_k^T) \mathcal{Z}(k) = \int_{\alpha_1}^{m_{\alpha_2}} H(X_z^\infty) \mathbf{f}(z) dz. \quad (4.24)$$

Proof: By the Portmanteau Theorem (cf. Theorem 2.1 in Billingsley [26]), to prove weak convergence, it suffices to demonstrate (4.24). Let \bar{H} denote an upper bound for $|H|$. It follows from Lemma 12, (4.23), and the fact that $|ab - cd| \leq |a+c||b-d| + |b+d||a-c|$ that for any $\epsilon > 0$, there exists $T_\epsilon < \infty$ s.t. for all $T \geq T_\epsilon$, and $k \in [\lfloor \alpha_1 T \rfloor + 1, m_{\alpha_2}(T)]$,

$$\left| H(X_k^T) (T \mathcal{Z}(k)) - H(X_{\frac{k}{T}}^\infty) \mathbf{f}(\frac{k}{T}) \right| \leq 2\bar{H}\epsilon + 4\bar{\mathbf{f}} \sup_{k \in [\lfloor \alpha_1 T \rfloor + 1, m_{\alpha_2}(T)]} |H(X_k^T) - H(X_{\frac{k}{T}}^\infty)|.$$

It follows from Lemma 12 that there similarly exists $T_{1,\epsilon} < \infty$ s.t. for all $T \geq T_{1,\epsilon}$, $\sup_{k \in [\lfloor \alpha_1 T \rfloor + 1, m_{\alpha_2}(T)]} |X_k^T - X_{\frac{k}{T}}^\infty| < \epsilon$. Combining the above with the fact that H is uniformly continuous implies that for any $\epsilon > 0$, there exists $T_{2,\epsilon} < \infty$ s.t. for all $T \geq T_{2,\epsilon}$,

$\sup_{k \in \llbracket \alpha_1 T \rrbracket + 1, m_{\alpha_2}(T)} |H(X_k^T)(T\mathcal{Z}(k)) - H(X_{\frac{k}{T}}^\infty)\mathbf{f}(\frac{k}{T})| < \epsilon$. Thus for any $\epsilon > 0$, there exists $T_{3,\epsilon} < \infty$ s.t. for all $T \geq T_{3,\epsilon}$, $\left| T^{-1} \sum_{k=\llbracket \alpha_1 T \rrbracket + 1}^{m_{\alpha_2}(T)} H(X_k^T)(T\mathcal{Z}(k)) - T^{-1} \sum_{k=\llbracket \alpha_1 T \rrbracket + 1}^{m_{\alpha_2}(T)} H(X_{\frac{k}{T}}^\infty)\mathbf{f}(\frac{k}{T}) \right| < \epsilon$. Noting that Riemann integrability implies $\lim_{T \rightarrow \infty} T^{-1} \sum_{k=\llbracket \alpha_1 T \rrbracket + 1}^{m_{\alpha_2}(T)} H(X_{\frac{k}{T}}^\infty)\mathbf{f}(\frac{k}{T}) = \int_{\alpha_1}^{m_{\alpha_2}} H(X_z^\infty)\mathbf{f}(z)dz$ completes the proof. \square

We now complete the proof of Theorem 4, and proceed by the standard path of demonstrating appropriate notions of tightness and convergence of finite-dimensional distributions. We begin by providing some additional background regarding the space $D([0, 1], \mathbb{R}^k)$ and associated notions of weak convergence. Recall that the space $D([0, 1], \mathbb{R}^k)$ contains all functions $x : [0, 1] \rightarrow \mathbb{R}^k$ that are right-continuous in $[0, 1)$ and have left-limits in $(0, 1]$ (i.e. cadlag on $[0, 1]$). Given such a cadlag function $x = (x_1, \dots, x_k) : [0, 1] \rightarrow \mathbb{R}^k$, let $\|x_j\| \triangleq \sup_{t \in [0, 1]} |x_j(t)|$, and $\|x\| \triangleq \sum_{j=1}^k \|x_j\|$. Given a set $\{\alpha_i\}_{i=0}^m$ satisfying $0 = \alpha_0 < \alpha_1 < \dots < \alpha_m = 1$, we say that $\{\alpha_i\}_{i=0}^m$ is δ -sparse if $\min_{1 \leq i \leq m} (\alpha_i - \alpha_{i-1}) \geq \delta$. Let $S_{\delta, m}$ denote the collection of all such δ -sparse sets consisting of exactly m points. Given x which is cadlag on $[0, 1]$ and $\delta \in (0, 1)$, let

$$w_x(\delta) \triangleq \inf_{m \in \mathbb{Z}^+, \{\alpha_i\} \in S_{\delta, m}} \max_{1 \leq i \leq m} \left(\sum_{j=1}^k \sup_{\alpha, \alpha' \in [\alpha_{i-1}, \alpha_i]} |x_j(\alpha) - x_j(\alpha')| \right),$$

where the infimum extends over all δ -sparse sets $\{\alpha_i\}$ of all possible finite lengths (i.e. values of m), where we note that one must only consider values of m which are at most $\lceil \delta^{-1} \rceil + 1$. Let us recall the following necessary and sufficient conditions for tightness in the space $D([0, 1], \mathbb{R}^k)$.

LEMMA 14 (Theorem 13.2, Billingsley [26]). *A sequence of probability measures $\{P_T\}_{T \geq 1}$ on the space $D([0, 1], \mathbb{R}^k)$ under the J_1 topology is tight if and only if the following conditions hold:*

$$\lim_{a \rightarrow \infty} \limsup_{T \rightarrow \infty} P_T(\|x\| \geq a) = 0; \quad (4.25)$$

$$\text{For all } \epsilon > 0, \quad \lim_{\delta \rightarrow 0} \limsup_{T \rightarrow \infty} P_T(w_x(\delta) \geq \epsilon) = 0. \quad (4.26)$$

We now use Lemma 14 to prove tightness of the sequence $\{\mathcal{M}^T(\alpha)_{0 \leq \alpha \leq 1}, T \geq 1\}$, by verifying (4.25) and (4.26).

LEMMA 15. *For $\mu \in (0, U)$, $\{\mathcal{M}^T(\alpha)_{0 \leq \alpha \leq 1}, T \geq 1\}$ is tight in $D([0, 1], \mathbb{R}^2)$.*

Proof: (4.25) follows immediately from the fact that $\mathcal{M}^T(\alpha)$ is supported on $[0, U]^2$ w.p.1 for all $\alpha \in [0, 1]$, $T \geq 1$. Let us prove (4.26). Recall from Corollary 2 that we may construct r.v.s Z^T, Y^T on the same probability space as \mathcal{M}^T s.t. the properties outlined in Corollary 2 hold. It follows from Lemma 12, and the Riemann integrability of \mathbf{f} , that there exists $\delta_0 > 0$ s.t. for all $\delta \in (0, \delta_0)$, and $\delta' \in (0, \delta)$, there exists $T_{\delta'}$ s.t. for all $T \geq T_{\delta'}$,

$$\sup_{t \in [1, \Lambda^T]} |X_t^T - X_{\frac{t}{T}}^\infty| < \delta' \quad , \quad \sup_{t \in [1, \Lambda^T]} |D_t^T - D_{\frac{t}{T}}^\infty| < \delta' \quad , \quad \left| \mathbb{P}\left(\frac{Z^T}{T} \leq \delta\right) - \int_0^\delta \mathbf{f}(z) dz \right| < \delta'. \quad (4.27)$$

Note that $X^\infty, D^\infty, \mathbf{f}$ are continuously differentiable on $[0, \Lambda^\infty + \epsilon_0]$ for some $\epsilon_0 > 0$. Let $\bar{D} \triangleq 32(1 + \sup_{0 \leq x \leq \Lambda^\infty + \epsilon_0} \max(|\frac{d}{dx} D_x^\infty|, |\frac{d}{dx} X_x^\infty|, |\frac{d}{dx} \mathbf{f}|))$. Also note that Lemma 12 implies that $\{Y^T, T \geq 1\}$ converges weakly to the r.v. which equals U w.p.1.

It thus follows from (4.27) and the triangle inequality that for all $\delta \in (0, \delta_0)$, there exists T_δ s.t. for all $T \geq T_\delta$,

$$\sup_{\substack{0 \leq \alpha_1 \leq \alpha_2 \leq 1 \\ \llbracket \alpha_2 T \rrbracket \leq \Lambda^T \\ |\alpha_1 - \alpha_2| \leq 2\delta}} \max \left(|X_{\llbracket \alpha_2 T \rrbracket}^T - X_{\llbracket \alpha_1 T \rrbracket}^T|, |D_{\llbracket \alpha_2 T \rrbracket}^T - D_{\llbracket \alpha_1 T \rrbracket}^T| \right) \leq \bar{D}\delta \quad , \quad |X_{\Lambda^T}^T - U| \leq \bar{D}\delta, \quad (4.28)$$

$$\mathbb{P}(\{\min(\frac{Z^T}{T}, 1 - \frac{Z^T}{T}) \leq 4\delta\} \cup \{Y^T = 0\}) \leq \bar{D}\delta.$$

As we will be referring to various sample-path constructions in the remainder of the proof, for a given $\omega \in \Omega$ (on the corresponding probability space), we will denote the dependence of various r.v.s on ω by adding ω as a subscript or superscript where appropriate. Given $\delta \in (0, \delta_0)$ and $T \geq T_\delta$, let $\Omega_{\delta, T} \subseteq \Omega$ denote that subset of Ω s.t. $\omega \in \Omega_{\delta, T}$ iff $\min(\frac{Z^{T, \omega}}{T}, 1 - \frac{Z^{T, \omega}}{T}) > 4\delta$ and $Y^{T, \omega} = U$. Let us fix some $\epsilon > 0, \delta \in (0, \delta_0), T \geq T_\delta, \omega \in \Omega_{\delta, T}$. For the given parameters, we now define a particular δ -sparse mesh. Note that for the given ω , there exists k_ω s.t. $k_\omega \delta \leq \frac{Z^{T, \omega}}{T} < (k_\omega + 1)\delta$, where $0 < (k_\omega - 2)\delta < (k_\omega + 2)\delta < 1$. Also, let $n_{1, \omega} \triangleq \sup\{n : \frac{Z^{T, \omega}}{T} + n\delta < 1 - \delta\}$, and $n_\omega \triangleq k_\omega + n_{1, \omega} + 2$. Then we let $\{\alpha_i^\omega, i \in [1, n_\omega]\}$ denote the following δ -sparse mesh. Let $\alpha_1^\omega = 0$. For $k \in [1, k_\omega - 1]$, let $\alpha_k^\omega = k\delta$. Let $\alpha_{k_\omega}^\omega = \frac{Z^{T, \omega}}{T}$. For $k \in [k_\omega + 1, k_\omega + 1 + n_{1, \omega}]$, let $\alpha_k^\omega = \frac{Z^{T, \omega}}{T} + (k - k_\omega)\delta$. Finally, we let $\alpha_{n_\omega}^\omega = 1$. Note that

$$\frac{Z^{T, \omega}}{T} = \alpha_{k_\omega}^\omega = \inf\{\alpha \in [0, 1] : \lfloor \alpha T \rfloor \geq Z^{T, \omega}\}, \quad (4.29)$$

$$i \in [2, n_\omega + 1], \alpha, \alpha' \in [\alpha_{i-1}^\omega, \alpha_i^\omega) \text{ implies } |\alpha - \alpha'| \leq 2\delta.$$

We treat two cases. First, suppose $Z^{T, \omega} \leq \Lambda^T - 1$. In this case, Corollary 2 implies that

$$\mathcal{M}^{T, \omega}(\alpha) = \begin{cases} (X_{\lfloor \alpha T \rfloor}^T, D_{\lfloor \alpha T \rfloor}^T) & \text{if } \lfloor \alpha T \rfloor \leq Z^{T, \omega} - 1; \\ (X_{Z^{T, \omega}}^T, 0) & \text{if } \lfloor \alpha T \rfloor \geq Z^{T, \omega}. \end{cases}$$

It thus follows from (4.28) and (4.29) that

$$\begin{aligned} \max_{i \in [2, k_\omega]} \sum_{j=1}^2 \sup_{\alpha, \alpha' \in [\alpha_{i-1}, \alpha_i)} |\mathcal{M}_j^{T, \omega}(\alpha) - \mathcal{M}_j^{T, \omega}(\alpha')| &\leq 2\bar{D}\delta, \\ \max_{i \in [k_\omega + 1, n_\omega + 1]} \sum_{j=1}^2 \sup_{\alpha, \alpha' \in [\alpha_{i-1}, \alpha_i)} |\mathcal{M}_j^{T, \omega}(\alpha) - \mathcal{M}_j^{T, \omega}(\alpha')| &= 0, \end{aligned}$$

and thus in this case $w_x^\omega(\delta) \leq 2\bar{D}\delta$.

Alternatively, suppose $Z^{T, \omega} = \Lambda^T$. In this case, Corollary 2 implies that

$$\mathcal{M}^T(\alpha) = \begin{cases} (X_{\lfloor \alpha T \rfloor}^T, D_{\lfloor \alpha T \rfloor}^T) & \text{if } \lfloor \alpha T \rfloor \leq \Lambda^T - 1; \\ (X_{\Lambda^T}^T, U) & \text{if } \lfloor \alpha T \rfloor = \Lambda^T; \\ (U, U) & \text{if } \lfloor \alpha T \rfloor \geq \Lambda^T + 1. \end{cases}$$

It again follows from (4.28) and (4.29) that

$$\max_{i \in [2, k_\omega]} \sum_{j=1}^2 \sup_{\alpha, \alpha' \in [\alpha_{i-1}, \alpha_i)} |\mathcal{M}_j^{T, \omega}(\alpha) - \mathcal{M}_j^{T, \omega}(\alpha')| \leq 2\bar{D}\delta.$$

Now, note that for $i \in [k_\omega + 1, n_\omega + 1]$ and $\alpha, \alpha' \in [\alpha_{i-1}, \alpha_i)$, $|\mathcal{M}_2^{T, \omega}(\alpha) - \mathcal{M}_2^{T, \omega}(\alpha')| = 0$, and $|\mathcal{M}_1^{T, \omega}(\alpha) - \mathcal{M}_1^{T, \omega}(\alpha')| \leq |X_{\Lambda^T}^T - U| \leq \bar{D}\delta$, with the final inequality following from (4.28). Thus in this case, it again holds that $w_x^\omega(\delta) \leq 2\bar{D}\delta$.

To complete the proof of (4.26), we need only carefully combine all of the above. In particular, let us fix any $\epsilon > 0$. It follows from (4.28) that for all $\delta \in (0, \delta_0)$ and $T \geq T_\delta$, $\mathbb{P}(\Omega_{\delta, T}) \geq 1 - \bar{D}\delta$, $w_x^\omega(\delta) \leq 2\bar{D}\delta$ for all $\omega \in \Omega_{\delta, T}$, and thus $\mathbb{P}(w_x^\omega(\delta) > 2\bar{D}\delta) \leq \bar{D}\delta$. We conclude that for all $\delta < \min(\delta_0, \frac{\epsilon}{3\bar{D}})$ and $T \geq T_\delta$, $\mathbb{P}(w_x^\omega(\delta) \geq \epsilon) \leq \bar{D}\delta$, from which (4.26) follows after taking limits in the appropriate order. \square

Next, we prove the following weak convergence of finite-dimensional distributions.

LEMMA 16. For $\mu \in (0, U)$, and $0 = \alpha_1 < \dots < \alpha_{k_0} = \Lambda^\infty < \dots < \alpha_n = 1$, $(\mathcal{M}^T(\alpha_1), \dots, \mathcal{M}^T(\alpha_n))$ converges weakly to $(\mathcal{M}^\infty(\alpha_1), \dots, \mathcal{M}^\infty(\alpha_n))$.

Proof: By the Portmanteau theorem, it suffices to prove that for any bounded uniformly continuous function $H : [0, U]^{2n} \rightarrow \mathbb{R}$,

$$\lim_{T \rightarrow \infty} \mathbb{E} [H(\mathcal{M}^T(\alpha_1), \dots, \mathcal{M}^T(\alpha_n))] = \mathbb{E} [H(\mathcal{M}^\infty(\alpha_1), \dots, \mathcal{M}^\infty(\alpha_n))]. \quad (4.30)$$

As a notational convenience, we define for $i \in [1, n]$,

$$\begin{aligned} H^\infty(i, x, y) &\triangleq H(X_{\alpha_1}^\infty, D_{\alpha_1}^\infty, \dots, X_{\alpha_i}^\infty, D_{\alpha_i}^\infty, x, y, x, y, \dots, x, y), \\ H^T(i, x, y) &\triangleq H(X_{[\alpha_1 T]}^T, D_{[\alpha_1 T]}^T, \dots, X_{[\alpha_i T]}^T, D_{[\alpha_i T]}^T, x, y, x, y, \dots, x, y), \\ \hat{H}^T(i, x, y) &\triangleq H(X_{[\alpha_1 T]}^T, D_{[\alpha_1 T]}^T, \dots, X_{[\alpha_i T]}^T, D_{[\alpha_i T]}^T, x, y, y, y, \dots, y, y). \end{aligned}$$

Then it follows from the definition of $\mathcal{M}^\infty(\alpha)$ that

$$\mathbb{E} [H(\mathcal{M}^\infty(\alpha_1), \dots, \mathcal{M}^\infty(\alpha_n))] = \sum_{i=1}^{k_0-1} \int_{\alpha_i}^{\alpha_{i+1}} H^\infty(i, X_z^\infty, 0) \mathbf{f}(z) dz + \gamma H^\infty(k_0, U, U).$$

It follows from Lemma 12 that there exists T_0 s.t. for all $T \geq T_0$, $[\alpha_{k_0-1} T] < \Lambda^T - 1 < \Lambda^T < [\alpha_{k_0+1} T]$. From Corollary 2, for all $T \geq T_0$,

$$\mathbb{E} [H(\mathcal{M}^T(\alpha_1), \dots, \mathcal{M}^T(\alpha_n))] = W_1^T + W_2^T + W_3^T + W_4^T,$$

where

$$\begin{aligned} W_1^T &\triangleq \sum_{i=1}^{k_0-2} \sum_{j=[\alpha_i T]+1}^{[\alpha_{i+1} T]} H^T(i, X_j^T, 0) \mathbb{P}(Z^T = j) + \sum_{j=[\alpha_{k_0-1} T]+1}^{\min\{[\alpha_{k_0} T], \Lambda^T-1\}} H^T(k_0-1, X_j^T, 0) \mathbb{P}(Z^T = j), \\ W_2^T &\triangleq I([\alpha_{k_0} T] < \Lambda^T - 1) \sum_{j=[\alpha_{k_0} T]+1}^{\Lambda^T-1} H^T(k_0, X_j^T, 0) \mathbb{P}(Z^T = j), \\ W_3^T &\triangleq \left[I([\alpha_{k_0} T] = \Lambda^T) \hat{H}^T(k_0-1, X_{\Lambda^T}^T, U) + I([\alpha_{k_0} T] > \Lambda^T) H^T(k_0-1, U, U) \right. \\ &\quad \left. + I([\alpha_{k_0} T] < \Lambda^T) H^T(k_0, U, U) \right] \mathbb{P}(Z^T = \Lambda^T) \mathbb{P}(Y^T = U), \\ W_4^T &\triangleq \left[I([\alpha_{k_0} T] \geq \Lambda^T) H^T(k_0-1, X_{\Lambda^T}^T, 0) \right. \\ &\quad \left. + I([\alpha_{k_0} T] < \Lambda^T) H^T(k_0, X_{\Lambda^T}^T, 0) \right] \mathbb{P}(Z^T = \Lambda^T) \mathbb{P}(Y^T = 0). \end{aligned}$$

Lemma 13 implies that $\lim_{T \rightarrow \infty} W_1^T = \sum_{i=1}^{k_0-1} \int_{\alpha_i}^{\alpha_{i+1}} H^\infty(i, X_z^\infty, 0) \mathbf{f}(z) dz$, and it follows from Lemma 12 and the uniform continuity of H that $\lim_{T \rightarrow \infty} W_4^T = \gamma H^\infty(k_0, U, U)$, and $\lim_{T \rightarrow \infty} W_2^T = \lim_{T \rightarrow \infty} W_3^T = 0$. Combining the above completes the proof of (4.30). \square

Proof: [Proof of Theorem 4] Combining Lemmas 14, 15, and 16 completes the proof of Theorem 4. \square

5. Conclusion

In this paper, we formulated and solved a dynamic distributionally robust multi-stage newsvendor model which naturally unifies the analysis of several inventory models with demand forecasting, and enables the optimizer to incorporate past realizations of demand into the structure of the uncertainty set going forwards. We explicitly computed the minimax optimal policy (and associated worst-case distribution) in closed form. Our main proof technique involved a non-trivial induction, combining ideas from convex analysis, probability, and DP. By analyzing in-depth the interplay between the minimax optimal policy and associated worst-case distribution, we proved that at optimality the worst-case demand distribution corresponds to the setting in which inventory may become obsolete at a random time, a scenario of practical interest. To gain further insight into our explicit solution, we computed the limiting dynamics as the time horizon grows large, by proving weak convergence to an appropriate limiting stochastic process, for which many quantities have a simple and intuitive closed form. We also compared to the analogous setting in which demand is independent across periods, and made several comparisons between these two models.

Our work leaves many interesting directions for future research. First, to make the model more practically applicable, it would be interesting to extend our results to more complicated dependency structures which robustify a larger class of forecasting models, e.g. that for which $E[D_{t+1}|D_{[t]}] = a_t D_t + b_t$ (i.e. general time-dependent linear conditional expectations). One could also incorporate higher order dependencies, non-linearities, as well as higher-dimensional covariate and feature spaces such as those arising in factor models. Similarly, one could consider inventory models with more realistic features, e.g. various notions of ordering costs, settings in which demand is lost, positive lead times, etc. Of course, our general approach could also be applied to many models used by the Operations Research community where forecasting and model uncertainty non-trivially interact. In all of these cases, it would be interesting to develop a more general theory of when the associated DP have a special structure which allows for a closed-form (or at least computationally tractable) solution. We note that such questions may also shed further light on the search for optimal probability inequalities in a variety of settings, e.g. bounds for functionals of martingale sequences and their generalizations.

Taking a broader view, perhaps the most interesting directions for future research involve exploring the relationship between different ways to model uncertainty in multi-stage optimization problems. What is the precise relationship between static and dynamic models of uncertainty, and how should we think about comparing and selecting between different models? This question involves notions of computational efficiency, model appropriateness in any given application, as well as questions related to the so-called price of correlations. Of course, such issues are intimately related to various relevant questions of a statistical nature. How precisely should our ability to fit and tune a model relate to our choice of uncertainty set? How does the notion of model over-fitting come into play here? Such questions become especially interesting in the context of multi-stage problems. Another intriguing line of questioning involves the relation between classical probability and (distributionally) robust optimization. For example, the obsolescence feature of the worst-case martingale in our setting can be interpreted as a “robust” manifestation of the celebrated martingale convergence theorem. A deeper understanding of the precise connection between the limit laws of classical probability and the solutions to distributionally robust optimization problems remains an intriguing open question, where we note that closely related questions have been the subject of several recent investigations (cf. Abernethy et al. [1], Bandi et al. [7, 8]). Formalizing and analyzing these and related questions will likely require a combination of ideas and techniques from optimization, statistics, and probability, all of which bring to bear different perspectives with which to understand uncertainty.

Acknowledgments

The authors would like to thank Erhan Bayraktar, Vishal Gupta, Garud Iyengar, Sandeep Juneja, Anton Kleywegt, Marcel Nutz, Alex Shapiro, and Alejandro Toriello for helpful discussions and insights. David A. Goldberg also gratefully acknowledges support from NSF grant no. 1453929.

References

- [1] Abernethy, J., P. Bartlett, R. Frongillo, A. Wibisono. 2013. How to hedge an option against an adversary: Black-scholes pricing is minimax optimal. *Advances in Neural Information Processing Systems* 2346–2354.
- [2] Agrawal, S., Y. Ding, A. Saberi, Y. Ye. 2012. Price of correlations in stochastic optimization. *Operations Research*. **60**(1) 150–162.
- [3] Aharon, B., G. Boaz, S. Shimrit. 2009. Robust multi-echelon multi-period inventory control. *European Journal of Operational Research*. **199**(3) 922–935.
- [4] Ahmed, S., U. Cakmak, A. Shapiro. 2007. Coherent risk measures in inventory problems. *European Journal of Operational Research*. **182** 226–238.
- [5] Altman, E., E. Feinberg, A. Shwartz. 2000. Weighted discounted stochastic games with perfect information. *Ann. Internat. Soc. Dynamic Games*. **5** 303–323.
- [6] R. Badinelli. 1990. The inventory costs of common mis-specification of demand-forecasting models. *The international journal of production research*. **28**(12) 2321–2340.
- [7] Bandi, C., D. Bertsimas, N. Youssef. 2015. Robust queueing theory. *Operations Research*.
- [8] Bandi, C., D. Bertsimas, N. Youssef. 2015. Robust transient multi-server queues and feed-forward networks. *Preprint*.
- [9] Bayraktar, E., Y. Zhang, Z. Zhou. 2013. A note on the Fundamental Theorem of Asset Pricing under model uncertainty. arXiv preprint arXiv:1309.2728.
- [10] Bayraktar, E., S. Yao. 2014. On the Robust Optimal Stopping Problem. *SIAM Journal on Control and Optimization*. **52**(5) 3135–3175.
- [11] Beigblock, M., M. Nutz. 2014. Martingale Inequalities and Deterministic Counterparts. *Electronic Journal of Probability*. **19**(95) 1–15.
- [12] Beigblock, M., M. Nutz, and N. Touzi. 2015. Complete Duality for Martingale Optimal Transport on the Line. arXiv preprint arXiv:1507.00671.
- [13] Bielecki, T., I. Cialenco, M. Pitera. 2014. On time consistency of dynamic risk and performance measures in discrete time. arXiv preprint arXiv:1409.7028.
- [14] Ben-Tal, A., L. El Ghaoui, A. Nemirovski. 2009. *Robust Optimization*. Princeton University Press, Princeton, NJ.
- [15] Ben-Tal, A., D. den Hertog, A. De Waegenaere, B. Melenberg, G. Rennen. 2013. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*. **59**(2) 341–357.
- [16] Ben-Tal, A., B. Golany, A. Nemirovski, J.P. Vial. 2005. Retailer-supplier flexible commitments contracts: a robust optimization approach. *Manufacturing & Service Operations Management*. **7** 248–271.
- [17] Bertsimas, D., X.V. Doan, K. Natarajan, C.P. Teo. 2010. Models for minimax stochastic linear optimization problems with risk aversion. *Mathematics of Operations Research*. **35**(3) 580–602.
- [18] Bertsimas, D., D. Brown, C. Caramanis. 2011. Theory and applications of robust optimization. *SIAM Review* **53**(3) 464–501.
- [19] Bertsimas, D., V. Gupta, N. Kallus. 2014. Robust SAA. arXiv preprint arXiv:1408.4445.
- [20] Bertsimas, D., N. Kallus. 2014. From Predictive to Prescriptive Analytics. arXiv preprint arXiv:1402.5481.
- [21] Bertsimas, D., D. A. Iancu, P.A. Parrilo. 2010. Optimality of affine policies in multistage robust optimization. *Mathematics of Operations Research*. **35**(2) 363–394.
- [22] Bertsimas, D., I. Popescu. 2005. Optimal inequalities in probability theory: A convex optimization approach. *SIAM Journal on Optimization*. **15**(3) 780–804.

-
- [23] Bertsimas, D., M. Sim, M. Zhang. 2014. A practicable framework for distributionally robust linear optimization. Preprint.
- [24] Bertsimas, D., A. Thiele. 2006. A robust optimization approach to inventory theory. *Operations Research*. **54**(1) 150-168.
- [25] Bertsimas, D. and Vayanos, P. 2015. Data-driven learning in dynamic pricing using adaptive optimization. Preprint.
- [26] Billingsley, P. 1999. *Convergence of Probability Measures*. Wiley, New York.
- [27] Blinder, A. 1983. Inventories and sticky prices: More on the microfoundations of macroeconomics. *National Bureau of Economic Research*.
- [28] Bonnans, J., A. Shapiro. 2013. *Perturbation analysis of optimization problems*. Springer Science and Business Media.
- [29] Box, G., G. Jenkins, G. Reinsel. 2011. *Time series analysis: forecasting and control*. Vol. 734. John Wiley and Sons.
- [30] Brown, G., J. Lu, R. Wolfson. 1964. Dynamic modeling of inventories subject to obsolescence. *Management Science*. **11**(1). 51-63.
- [31] Carrizosaa, E., A. Olivares-Nadal, P. Ramirez-Cobob. 2014. Robust newsvendor problem with autoregressive demand. *Preprint*.
- [32] Chen, F., Z. Drezner, J. Ryan, D. Simchi-Levi. 2000. Quantifying the bullwhip effect in a simple supply chain: The impact of forecasting, lead times, and information. *Management Science*. **46**(3) 436-443.
- [33] Chen, Y., G. Iyengar. 2011. A new robust cycle-based inventory control policy. working paper.
- [34] Cheng, X., F. Riedel. 2013. Optimal stopping under ambiguity in continuous time. *Mathematics and Financial Economics*. **7**(1) 29-68.
- [35] Choi, S., A. Ruszczyński. 2008. A risk-averse newsvendor with law invariant coherent measures of risk. *Operations Research Letters*. **36**(1) 77-82.
- [36] Clark, A. J., H. Scarf. 1960. Optimal policies for a multi-echelon inventory problem. *Management Science*. **50**(12) 475-490.
- [37] Cerag, P. 2010. Advances in Inventory Management: Dynamic Models. *Erasmus Research Institute of Management (ERIM)*. No. EPS-2010-199-LIS.
- [38] Chao, X., X. Gong, C. Shi, H. Zhang. 2015. Approximation Algorithms for Perishable Inventory Systems. *Operations Research*. **63**(3) 585-601.
- [39] Cogger, K. 1988. Proposals for research in time series forecasting. *International Journal of Forecasting*. **4**(3) 403-410.
- [40] Delage, E., Y. Ye. 2010. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research*. **58**(3) 596-612.
- [41] Dolinsky, Y., Soner, M. 2014. Martingale optimal transport and robust hedging in continuous time. *Probability Theory and Related Fields*. **160**(1-2) 391-427.
- [42] Dominguez, M., I. Lobato. 2004. Consistent estimation of models defined by conditional moment restrictions. *Econometrica*. **72**(5) 1601-1615.
- [43] Dupacová, J. 1987. The minimax approach to stochastic programming and an illustrative application. *Stochastics*. **20** 73-88.
- [44] Dupacová, J. 2001. Stochastic programming: Minimax approach. In: A. Floudas and P.M. Pardalos (Eds.), *Encyclopedia of Optimization*. **5** 327-330. Kluwer Academic Publishers.
- [45] Edgeworth, F. 1888. The mathematical theory of banking. *Journal of the Royal Statistical Society*. **51**(1) 113-127.
- [46] Epstein, L. G., M. Schneider. 2003. Recursive multiple-priors. *Journal of Economic Theory*. **113**(1) 1-31.
- [47] Erkip, N., W. Hausman, S. Nahmias. 1990. Optimal centralized ordering policies in multi-echelon inventory systems with correlated demands. *Management Science*. **36**(3) 381-392.

- [48] Escanciano, J., I. Lobato. 2009. Testing the martingale hypothesis. Palgrave Hand-book of Econometrics, Palgrave, MacMillan.
- [49] Feldman, R. 1978. A continuous review (s, S) inventory system in a random environment. *Journal of Applied Probability*. 654–659.
- [50] Fildes, R., S. Makridakis. 1995. The impact of empirical accuracy studies on time series analysis and forecasting. *International Statistical Review*. 289-308.
- [51] Fildes, R., K. Nikolopoulos, S. F. Crone, A. A. Syntetos. 2008. Forecasting and Operational Research: A Review. *The Journal of the Operational Research Society*. **59**(9) 1150–1172.
- [52] Föllmer, H., P. Leukert. 1999. Quantile hedging. *Finance and Stochastics*. **3**(3) 251-273.
- [53] Franke, J. 1985. Minimax-robust prediction of discrete time series. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*. **68**(3) 337-364.
- [54] Gallego, G. 1998. New bounds and heuristics for (Q, r) policies. *Management Science*. **44**(2) 219–233.
- [55] Gallego, G., O. Ozer. 2001. Integrating replenishment decisions with advance demand information. *Management Science*. **47**(10) 1344–1360.
- [56] Gallego, G. 2001. Minimax analysis for finite-horizon inventory models. *IIE Transactions*. **33**(10) 861–874.
- [57] Gallego, G., I. Moon. 1993. The distribution free newsboy problem: review and extensions. *Journal of the Operational Research Society*. **44**(8) 825–834.
- [58] Gallego, G., I. Moon. 1994. Distribution free procedures for some inventory models. *Journal of the Operational research Society*. 651–658.
- [59] Gallego, G., J. Ryan, D. Simchi-Levi. 2001. Minimax analysis for finite horizon inventory models. *IIE Transactions*. **33** 861–874.
- [60] Glasserman, P., X. Xu. 2014. Robust risk measurement and model risk. *Quantitative Finance*. **14**(1) 29–58.
- [61] De Gooijer, J., R. Hyndman. 2006. 25 years of time series forecasting. *International journal of forecasting*. **22**(3) 443–473.
- [62] Graves, S.C. 1999. A single-item inventory model for a nonstationary demand process. *Manufacturing & Service Operations Management*. **1**(1) 50–61.
- [63] Graves, S.C., M.C. Meal, S. Dasu, Y. Qiu. 1986. Two-stage production planning in a dynamic environment. Axsater S, Schneeweiss C, Silver E, eds. *Multi-Stage Production Planning and Inventory Control*. Lecture Notes in Economics and Mathematical Systems. Vol. 266. Springer-Verlag, Berlin. 9-43.
- [64] Guerrier, S., R. Molinari, M. Victoria-Feser. 2014. Estimation of Time Series Models via Robust Wavelet Variance. *Austrian Journal of Statistics*. **43**(4) 267–277.
- [65] Gupta, Vishal. 2015. Near-Optimal Ambiguity Sets for Distributionally Robust Optimization. Preprint.
- [66] Hanasusanto, G., A. Grani, D. Kuhn, W. Wallace, S. Zymler. 2012. Distributionally robust multi-item newsvendor problems with multimodal demand distributions. *Mathematical Programming*. 1–32.
- [67] Hansen, L., T. Sargent. 2001. Robust control and model uncertainty. *American Economic Review*. 60–66.
- [68] Hausman, W.H. 1969. Sequential decision problems: A model to exploit existing forecasters. *Management Science*. **16**(2) B-93-B-111.
- [69] Heath, D. C., P. L. Jackson. 1994. Modeling the evolution of demand forecasts with application. *IIE Transactions*. **26**(3) 17-30.
- [70] Hill, T., R. Kertz. 1983. Stop rule inequalities for uniformly bounded sequences of random variables. *Transactions of the American Mathematical Society*. **278** 197–207.
- [71] Huber, P. 1964. Robust estimation of a location parameter. *The Annals of Mathematical Statistics*. **35**(1) 73–101.
- [72] Iancu, D.A., M. Petrik, D. Subramanian. 2013. Tight Approximations of Dynamic Risk Measures. *Mathematics of Operations Research*, to appear.

-
- [73] Iancu, D.A., M. Sharma, M. Sviridenko. 2013. Supermodularity and affine policies in dynamic robust optimization. *Operations Research*. **61**(4) 941-956.
- [74] Iglehart, D. L. 1963. Optimality of (s, S) policies in the infinite horizon dynamic inventory problem. *Management Science*. **9**(2) 259-267.
- [75] Iglehart, D., S. Karlin. 1962. Optimal policy for dynamic inventory process with nonstationary stochastic demands. **8**.
- [76] Iida, T., P. Zipkin. 2006. Approximate solutions of a dynamic forecasting inventory model. *Manufacturing & Service Operations Management*. **8**(4) 407-425.
- [77] Ishihara, J., M. Terra, J. Campos. 2006. Robust Kalman filter for descriptor systems. *IEEE Transactions on Automatic Control*. **51**(8) 1354.
- [78] Iyengar, G. 2005. Robust dynamic programming. *Mathematics of Operations Research*. **30**(2) 257-280.
- [79] Joglekar, P., P. Lee. 1993. An exact formulation of inventory costs and optimal lot size in face of sudden obsolescence. *Operations Research Letters*. **14**(5) 283-290.
- [80] Johnson, G. D., H. E. Thompson. 1975. Optimality of myopic inventory policies for certain dependent demand processes. *Management Science*. **21**(11) 1303-1307.
- [81] Kalymon, B. 1971. Stochastic prices in a single-item inventory purchasing model. *Operations Research*. **19**(6) 1434-1458.
- [82] Kasugai, H., T. Kasegai. 1961. Note on minimax regret ordering policy-static and dynamic solutions and a comparison to maximin policy. *Journal of the Operations Research Society of Japan*. **3**(4) 155-169.
- [83] Klabjan, D., D. Simchi-Levi, M. Song. 2013. Robust stochastic lot-sizing by means of histograms. *Production and Operations Management*. **22** 691-710.
- [84] Lam, H. 2011. Sensitivity to Serial Dependency of Input Processes: A Robust Approach. Preprint.
- [85] Lee, H., V. Padmanabhan, S. Whang. 2004. Information distortion in a supply chain: the bullwhip effect. *Management Science*. **50**(12) 1875-1886.
- [86] Levi, R., G. Janakiraman, M. Nagarajan. 2008. A 2-approximation algorithm for stochastic inventory control models with lost sales. *Mathematics of Operations Research*. **33**(2) 351-374.
- [87] Lovejoy, W. 1992. Stopped myopic policies for some inventory models with uncertain demand distributions. *Management Science*. **38**(5) 688-707.
- [88] Lu, X. W., J. S. Song, A. Regan. 2006. Inventory planning with forecast updates: Approximate solutions and cost error bounds. *Operations Research*. **54**(6) 1079-1097.
- [89] Luz, M., M. Moklyachuk. 2014. Minimax-robust filtering problem for stochastic sequences with stationary increments. *Theory of Probability and Mathematical Statistics*. **89** 127-142.
- [90] Maronna, R., M. Douglas, V. Yohai. 2006. Robust statistics. John Wiley and Sons, Chichester.
- [91] Marandi, A., D. Hertog. 2015. When are static and adjustable robust optimization with constraint-wise uncertainty equivalent?. Preprint.
- [92] Medina, M., E. Ronchetti. 2015. Robust statistics: a selective overview and new directions. *Wiley Interdisciplinary Reviews: Computational Statistics*.
- [93] Miller, B. 1986. Scarf's state reduction method, flexibility, and a dependent demand inventory model. *Operations Research*. **34**(1) 83-90.
- [94] Milner, J. M., P. Kouvelis. 2005. Order quantity and timing flexibility in supply chains: The role of demand characteristics. *Management Science*. **51**(6) 970-985.
- [95] Moklyachuk, M. 2000. Robust procedures in time series analysis. *Theory Stoch. Process*. **6**(22) 3-4.
- [96] Nilim, A., L. El Ghaoui. 2005. Robust control of Markov decision processes with uncertain transition matrices. *Operations Research*. **53**(5) 780-798.
- [97] Nishimura, K.G., H. Ozaki. 2007. Irreversible investment and Knightian uncertainty. *Journal of Economic Theory*. **136** 668-694.
- [98] Notzon, E. 1970. Minimax inventory and queueing models. Technical Report No. 7, Department of Operations Research, Stanford University.

- [99] Perakis, G., G. Roels. 2008. Regret in the newsvendor model with partial information. *Operations Research*. **56**(1) 188-203.
- [100] Pindyck, R. 1982. Adjustment costs, uncertainty, and the behavior of the firm. *The American Economic Review*. 415–427.
- [101] Pierskalla, W. 1969. An inventory problem with obsolescence. *Naval Research Logistics Quarterly*. **16**(2). 217–228.
- [102] Popescu, I. 2007. Robust mean-covariance solutions for stochastic optimization. *Operations Research*. **55**(4) 98-112.
- [103] Prékopa, A. 1995. *Stochastic Programming*. Kluwer Academic Publishers.
- [104] Puterman, Martin L. 2014. Markov decision processes: discrete stochastic dynamic programming. John Wiley and Sons.
- [105] Reyman, G. 1989. State reduction in a dependent demand inventory model given by a time series. *European Journal of Operational Research*. **41**(2) 174-180.
- [106] Riedel, F. 2009. Optimal stopping with multiple priors. *Econometrica*. **77**(3) 857-908.
- [107] Royden, H., P. Fitzpatrick. 1988. *Real analysis*. New York: Prentice Hall.
- [108] Ruckdeschel, P. 2010. Optimally (distributional-) robust Kalman filtering. arXiv preprint arXiv:1004.3393.
- [109] Ryan, J. 1998. Analysis of inventory models with limited demand information. Ph.D. Thesis.
- [110] Scarf, H. 1958. A min-max solution of an inventory problem. In K. Arrow, S. Karlin and H. Scarf (Eds.) *Studies in the Mathematical Theory of Inventory and Production*. Stanford University Press, Stanford, CA, 201–209.
- [111] Scarf, H. 1959. Bayes solution of the statistical inventory problem. *The annals of mathematical statistics*. 490–508.
- [112] Scarf, H. 1960. Some remarks on Bayes solutions to the inventory problem. *Naval research logistics quarterly*. **7**(4) 591–596.
- [113] Scarf, H. 1960. The Optimality of (s, S) Policies in the Dynamic Inventory Problem. *Mathematical Methods in the Social Sciences*. Stanford University Press. 196-202.
- [114] See, C. T., M. Sim. 2010. Robust approximation to multiperiod inventory management. *Operations Research*. **58**(3) 583-594.
- [115] Shapiro, A. 2012. Minimax and risk averse multistage stochastic programming. *European Journal of Operational Research*. **219** 719–726.
- [116] Shapiro, A. 2009. On a time consistency concept in risk averse multistage stochastic programming. *Operations Research Letters*. **37**(3) 143–147.
- [117] Shapiro, A., D. Dentcheva, A. Ruszczyński. 2009. *Lectures on Stochastic Programming: Modeling and Theory*. SIAM, Philadelphia.
- [118] Shapiro, A. Rectangular sets of probability measures. 2015. Optimization online.
- [119] Song, J., P. Zipkin. 1993. Inventory control in a fluctuating demand environment. *Operations Research*. **41**(2) 351-370.
- [120] Song, J., P. Zipkin. 1996. Managing inventory with the prospect of obsolescence. *Operations Research*. **44**(1) 215–222.
- [121] Stockinger, N., R. Dutter. 1987. Robust time series analysis: A survey. *Kybernetika*. **23**(7) 3–88.
- [122] Toktay, L., L. Wein. 2001. Analysis of a forecasting-productioninventory system with stationary demand. *Management Science*. **47**(9) 1268-1281.
- [123] Trojanowska, M., P.M. Kort. 2010. The Worst Case for Real Options. *Journal of Optimization Theory and Applications*. **146**(3) 709-734.
- [124] Veinott, Jr., A. F. 1965. Optimal policy for a multi-product, dynamic, nonstationary inventory problem. *Management Science*. **12**(3) 206-222.

- [125] Tukey, W. 1960. A survey of sampling from contaminated distributions. *Contributions to probability and statistics*. **2** 448–485.
- [126] Veinott, Jr., A. F. 1966. On the optimality of (s, S) inventory policies: new conditions and a new proof. *SIAM Journal on Applied Mathematics*. **14** 1067–1083.
- [127] Wang, T., A. Atasu, M. Kurtulus. 2012. A multi-ordering newsvendor model with dynamic forecast evolution. *Manufacturing & Service Operations Management*. **14**(3) 472–484.
- [128] Whitt, W. 2002. *Stochastic-process limits: an introduction to stochastic-process limits and their application to queues*. Springer.
- [129] Wiesemann, W., D. Kuhn, B. Rustem. 2013. Robust Markov decision processes. *Mathematics of Operations Research*. **38**(1) 153–183.
- [130] Xin, L., D.A. Goldberg, and A. Shapiro. "Time (in) consistency of multistage distributionally robust inventory models with moment constraints." arXiv preprint arXiv:1304.3074 (2013).
- [131] Žáčková, J. 1966. On minimax solution of stochastic linear programming problems. *Časopis Pěstování Matematiky*. **91** 423–430.
- [132] Yue, J., B. Chen, M. Wang. 2006. Expected value of distribution information for the newsvendor problem. *Operations Research*. **54**(6) 1128–1136.
- [133] Zipkin, P.H. 2000. *Foundations of Inventory Management*. McGraw-Hill, Boston.
- [134] Zhu, Z., J. Zhang, Y. Ye. 2013. Newsvendor optimization with limited distribution information. *Optimization methods and software*. **28**(3) 640–667.

6. Technical Appendix

6.1. Proof of Lemma 1

Proof of Lemma 1: We first prove that one can w.l.o.g. restrict to policies always ordering up to at least 0. For $s \in [1, T]$, let $\hat{\Pi}^{s,0}$ denote that subset of $\hat{\Pi}$ consisting of those policies $\hat{\pi}$ with the following property: for all $t \in [T - s + 1, T]$ and $\mathbf{d} = (d_1, \dots, d_T) \in [0, U]_{\mathcal{Q}}^T$, $x_t^{\hat{\pi}}(\mathbf{d}_{[t-1]}) \geq 0$. Also, let $\hat{\Pi}^{0,0} \triangleq \hat{\Pi}$. We now prove that for all $s \in [0, T - 1]$, and any policy $\hat{\pi} \in \hat{\Pi}^{s,0}$, there exists a policy $\hat{\pi}' \in \hat{\Pi}^{s+1,0}$ s.t. for all vectors $\mathbf{d} = (d_1, \dots, d_T) \in [0, U]_{\mathcal{Q}}^T$ and $t \in [1, T]$, $C_t^{\hat{\pi}'}(\mathbf{d}_{[t]}) \leq C_t^{\hat{\pi}}(\mathbf{d}_{[t]})$, proving the desired result. Let us fix any $s \in [0, T - 1]$, and let $\hat{\pi}$ be any policy belonging to $\hat{\Pi}^{s,0}$. We now define a sequence of functions $\{x'_t, t \in [1, T]\}$, where x'_1 is a real constant, and x'_t is a map from $[0, U]_{\mathcal{Q}}^{t-1}$ to \mathcal{Q} for $t \geq 2$. We will then prove that these functions correspond to a policy belonging to $\hat{\Pi}^{s+1,0}$, which incurs cost no greater than $\hat{\pi}$ in every time period for every sample path of demand. Let $x'_t = x_t^{\hat{\pi}}$ for all $t \in [1, T] \setminus \{T - s\}$, and $x'_{T-s} = \max(0, x_{T-s}^{\hat{\pi}})$. We now verify that $\{x'_t, t \in [1, T]\}$ corresponds to some policy belonging to $\hat{\Pi}^{s+1,0}$, and begin by demonstrating that for all $t \in [0, T - 1]$ and $\mathbf{d} = (d_1, \dots, d_t) \in [0, U]_{\mathcal{Q}}^t$, $x'_{t+1}(\mathbf{d}) \geq x'_t(\mathbf{d}_{[t-1]}) - d_t$, where $x'_1(d_{[0]}) \triangleq x'_1$, $x'_0(\mathbf{d}_{[-1]}) \triangleq x_0$, and $d_0 \triangleq 0$. The assertion is immediate for $t \in [1, T] \setminus \{T - s\}$, as in these cases the property is inherited from $\{x_t^{\hat{\pi}}, t \in [1, T]\}$, and the monotonicity of our construction. We now verify the assertion for $t = T - s$, by demonstrating that $x'_{T-s+1}(\mathbf{d}) \geq x'_{T-s}(\mathbf{d}_{[T-s-1]}) - d_{T-s}$, and proceed by a case analysis. First, suppose $x_{T-s}^{\hat{\pi}}(\mathbf{d}_{[T-s-1]}) \geq 0$. In this case, the desired property is trivially inherited from $x_{T-s}^{\hat{\pi}}$. If not, by construction $x'_{T-s}(\mathbf{d}_{[T-s-1]}) = 0$, and thus $x'_{T-s}(\mathbf{d}_{[T-s-1]}) - d_{T-s} \leq 0$. However, $\hat{\pi} \in \hat{\Pi}^{s,0}$ implies that $x_{T-s+1}^{\hat{\pi}}(\mathbf{d}) \geq 0$, demonstrating the desired property. Combining the above with the fact that $x'_t \geq 0$ for all $t \in [T - s + 1, T]$ by virtue of x'_t equalling $x_t^{\hat{\pi}}$ for all $t \geq T - s + 1$ and $\hat{\pi}$ belonging to $\hat{\Pi}^{s,0}$, and the fact that $x'_{T-s} \geq 0$ by construction, completes the proof that for all $s \in [0, T - 1]$, and any policy $\hat{\pi} \in \hat{\Pi}^{s,0}$, there exists a policy $\hat{\pi}' \in \hat{\Pi}^{s+1,0}$ s.t. $x_t^{\hat{\pi}'} = x_t^{\hat{\pi}}$ for all $t \in [1, T] \setminus \{T - s\}$, and $x_{T-s}^{\hat{\pi}'} = \max(0, x_{T-s}^{\hat{\pi}})$. Combining with the fact that for all $d \in [0, U]$ and $x < 0$,

$$C(x, d) = b(d + |x|) > b[d - 0]_+ + [0 - d]_+ = bd,$$

further demonstrates that for all vectors $\mathbf{d} = (d_1, \dots, d_T) \in [0, U]_{\mathcal{Q}}^T$ and $t \in [1, T]$, $C_t^{\hat{\pi}'}(\mathbf{d}_{[t]}) \leq C_t^{\hat{\pi}}(\mathbf{d}_{[t]})$, completing the proof of the desired statement. Applying the statement inductively, we conclude

that for any $\hat{\pi} \in \hat{\Pi}$, there exists a policy $\hat{\pi}' \in \hat{\Pi}^{T,0}$ (i.e. a policy ordering up to at least 0 in every period for all sample paths) which incurs cost at most that incurred by $\hat{\pi}$.

We next prove that one can restrict to policies always ordering up to at most U . For $s \in [1, T]$, let $\hat{\Pi}^{s,U}$ denote that subset of $\hat{\Pi}^{T,0}$ consisting of those $\hat{\pi}$ with the following property: for all $t \in [T-s+1, T]$ and $\mathbf{d} = (d_1, \dots, d_T) \in [0, U]_{\mathcal{Q}}^T$, $x_t^{\hat{\pi}}(\mathbf{d}_{[t-1]}) \leq U$ whenever $y_t^{\hat{\pi}}(\mathbf{d}_{[t-1]}) \leq U$. Also, let $\hat{\Pi}^{0,U} \triangleq \hat{\Pi}^{T,0}$. We now prove that for all $s \in [0, T-1]$, and any policy $\hat{\pi} \in \hat{\Pi}^{s,U}$, there exists a policy $\hat{\pi}' \in \hat{\Pi}^{s+1,U}$ s.t. for all vectors $\mathbf{d} = (d_1, \dots, d_T) \in [0, U]_{\mathcal{Q}}^T$ and $t \in [1, T]$, $C_t^{\hat{\pi}'}(\mathbf{d}_{[t]}) \leq C_t^{\hat{\pi}}(\mathbf{d}_{[t]})$. Let us fix any $s \in [0, T-1]$, and let $\hat{\pi}$ be any policy belonging to $\hat{\Pi}^{s,U}$. We now define a sequence of functions $\{x'_t, t \in [1, T]\}$, where x'_1 is a real constant, and x'_t is a map from $[0, U]_{\mathcal{Q}}^{t-1}$ to \mathcal{Q} for $t \geq 2$. We will then prove that these functions correspond to a policy belonging to $\hat{\Pi}^{s+1,U}$, which incurs cost no greater than $\hat{\pi}$ in every time period for every sample path of demand. Let $x'_t = x_t^{\hat{\pi}}$ for all $t \in [1, T] \setminus \{T-s\}$; and for $\mathbf{d} \in [0, U]_{\mathcal{Q}}^{T-s-1}$ let

$$x'_{T-s}(\mathbf{d}) \triangleq \begin{cases} \min(U, x_{T-s}^{\hat{\pi}}(\mathbf{d})) & \text{if } y_{T-s}^{\hat{\pi}}(\mathbf{d}) \leq U, \\ x_{T-s}^{\hat{\pi}}(\mathbf{d}) & \text{else.} \end{cases}$$

We now verify that $\{x'_t, t \in [1, T]\}$ corresponds to some policy belonging to $\hat{\Pi}^{s+1,U}$, and begin by demonstrating that for all $t \in [0, T-1]$ and $\mathbf{d} = (d_1, \dots, d_t) \in [0, U]_{\mathcal{Q}}^t$, $x'_{t+1}(\mathbf{d}) \geq x'_t(\mathbf{d}_{[t-1]}) - d_t$. The assertion is immediate for $t \in [1, T] \setminus \{T-s-1\}$, as in these cases the property is inherited from $\{x_t^{\hat{\pi}}, t \in [1, T]\}$, and the monotonicity of our construction. We now verify the assertion for $t = T-s-1$, by demonstrating that $x'_{T-s}(\mathbf{d}) \geq x'_{T-s-1}(\mathbf{d}_{[T-s-1]}) - d_{T-s-1}$, and proceed by a case analysis. First, suppose that either $y_{T-s}^{\hat{\pi}}(\mathbf{d}_{[T-s-1]}) > U$, or $y_{T-s}^{\hat{\pi}}(\mathbf{d}_{[T-s-1]}) < U$ and $x_{T-s}^{\hat{\pi}}(\mathbf{d}_{[T-s]}) \leq U$. In this case, the desired property is either trivially true, or inherited from $x_{T-s}^{\hat{\pi}}$. If not, it must hold that $y_{T-s}^{\hat{\pi}}(\mathbf{d}_{[T-s]}) \leq U$, and $x_{T-s}^{\hat{\pi}}(\mathbf{d}_{[T-s]}) > U$. However, by construction this implies that $x'_{T-s-1}(\mathbf{d}_{[T-s-2]}) - d_{T-s-1} \leq U$, while $x'_{T-s}(\mathbf{d}_{[T-s-1]}) = U$, demonstrating the desired property. Combining the above with the fact that $x'_t(\mathbf{d}) \leq U$ whenever $x'_{t-1}(\mathbf{d}_{[t-2]})$ for all $\mathbf{d} \in [0, U]_{\mathcal{Q}}^{t-1}$ and $t \in [T-s+1, T]$ by virtue of x'_t equalling $x_t^{\hat{\pi}}$ for all $t \geq T-s+1$ and $\hat{\pi}$ belonging to $\hat{\Pi}^{s,U}$, and the fact that $x'_{T-s} \leq U$ by construction, completes the proof that for all $s \in [0, T-1]$, and any policy $\hat{\pi} \in \hat{\Pi}^{s,U}$, there exists a policy $\hat{\pi}' \in \hat{\Pi}^{s+1,U}$ s.t. $x_t^{\hat{\pi}'} = x_t^{\hat{\pi}}$ for all $t \in [1, T] \setminus \{T-s\}$, and x'_{T-s} is as defined above. Combining with the fact that for all $d \in [0, U]$ and $x > U$,

$$C(x, d) = x - d > b[d - U]_+ + [U - d]_+ = U - d,$$

further demonstrates that for all vectors $\mathbf{d} = (d_1, \dots, d_T) \in [0, U]_{\mathcal{Q}}^T$ and $t \in [1, T]$, $C_t^{\hat{\pi}'}(\mathbf{d}_{[t]}) \leq C_t^{\hat{\pi}}(\mathbf{d}_{[t]})$, completing the proof of the desired statement. Applying the statement inductively, we conclude that for any $\hat{\pi} \in \hat{\Pi}$, so long as $x_0 \leq U$, there exists a policy $\hat{\pi}' \in \hat{\Pi}^{T,U}$ (i.e. a policy ordering up to a level in $[0, U]$ in every period for all sample paths) which incurs cost at most that incurred by $\hat{\pi}$, completing the proof. \square .

6.2. Proof of Theorem 2

Our proof is quite similar, conceptually, to the proof in Iyengar [78] that a robust MDP satisfying the so-called recangularity property can be solved by dynamic programming. However, as our problem does not seem to fit precisely into that framework (e.g. since the conditional distribution of demand in a given time period can in principle depend on the entire history of demand up to that time), and as it is well-known that the rectangularity property and related notions can be quite subtle (cf. Xin, Goldberg and Shapiro [130]), we include a self-contained proof for completeness. We will proceed by a backwards induction, reasoning inductively that w.l.o.g. in period T the conditional distribution of demand can be chosen to be consistent with \hat{Q}^* , and thus the policy in

period T can be chosen to be consistent with $\hat{\pi}^*$, and thus the conditional distribution of demand in period $T-1$ can be chosen to be consistent with \hat{Q}^* , etc. It is worth noting that one could also proceed by using e.g. minmax theorems, strong duality, saddle-point theory, and/or general results from the theory of stochastic games, and we refer the interested reader to Nilim and El Ghaoui [96] for closely related results proved using such alternative approaches. We begin by making several definitions. For $\pi \in \Pi$, let $\mathbf{MAR}_0^\pi = \mathbf{MAR}_T$. For $t \in [1, T-1]$, let us define \mathbf{MAR}_t^π as follows. $Q \in \mathbf{MAR}_t^\pi$ iff $Q \in \mathbf{MAR}_{t-1}^\pi$, and for all $\mathbf{q} \in \text{supp}(Q_{[T-t]})$, $Q_{T-t+1|\mathbf{q}} = \bar{Q}_{x_{T-t+1}^\pi(\mathbf{q}), q_{T-t}}$. Also, let \mathbf{MAR}_T^π denote the singleton $\{\hat{Q}^*\}$. Similarly, let $\Pi_0 = \Pi$, and for $t \in [1, T-1]$, let us define Π_t as follows. $\pi \in \Pi_t$ iff $\pi \in \Pi_{t-1}$, and for all $\mathbf{q} \in [0, U]_{\mathcal{Q}}^{T-t}$, $x_{T-t+1}^\pi(\mathbf{q}) = \bar{\Phi}_{y_{T-t+1}^\pi(\mathbf{q}), q_{T-t}}$. Also, let Π_T denote the singleton $\{\hat{\pi}^*\}$.

We now state a lemma which follows immediately from our definitions, and formalizes the notion that for any probability measure $Q \in \mathbf{MAR}$, and time t , there is a unique probability measure belonging to \mathbf{MAR}_t^π which agrees with Q in periods $[1, t-1]$, and can be constructed naturally through an appropriate composition of marginal distributions; as well as the analogous statement for constructing policies belonging to Π_t which agree with a given policy in periods $[1, t-1]$. The associated measures (and policies) can in all cases be constructed through the natural composition of marginals (measurable functions).

LEMMA 17. *Given $\pi \in \Pi$, $Q \in \mathbf{MAR}_T$, and $t \in [1, T-1]$, there exists a unique probability measure $Q^{\pi, t} \in \mathbf{MAR}_t^\pi$ satisfying $Q_{[T-t]}^{\pi, t} = Q_{[T-t]}$. Similarly, given $\pi \in \Pi$ and $t \in [1, T-1]$, there exists a unique policy $\pi^t \in \Pi_t$ satisfying $x_{[T-t]}^{\pi^t} = x_{[T-t]}^\pi$.*

We now prove (by induction) a theorem which formalizes the intuitive backwards induction described above, and which will directly imply Theorem 2.

THEOREM 8. *For all $s \in [0, T-1]$, $\pi \in \Pi_s$, and $Q \in \mathbf{MAR}_s^\pi$,*

$$\mathbb{E}_Q \left[\sum_{t=T-s}^T C_t^\pi \right] = \sum_{\mathbf{q} \in \text{supp}(Q_{[T-s-1]})} \left(\sum_{z \in [0, U]_{\mathcal{Q}}} \hat{f}^{s+1}(x_{T-s}^\pi(\mathbf{q}), z) Q_{T-s|\mathbf{q}}(z) \right) Q_{[T-s-1]}(\mathbf{q}). \quad (6.31)$$

For all $s \in [0, T-1]$, $\pi \in \Pi_s$, and $Q \in \mathbf{MAR}_{s+1}^\pi$,

$$\mathbb{E}_Q \left[\sum_{t=T-s}^T C_t^\pi \right] = \sum_{\mathbf{q} \in \text{supp}(Q_{[T-s-1]})} \hat{g}^{s+1}(x_{T-s}^\pi(\mathbf{q}), q_{T-s-1}) Q_{[T-s-1]}(\mathbf{q}). \quad (6.32)$$

For all $s \in [0, T-1]$ and $\pi \in \Pi_s$,

$$\sup_{Q \in \mathbf{MAR}} \mathbb{E}_Q \left[\sum_{t=1}^T C_t^\pi \right] = \sup_{Q \in \mathbf{MAR}_{s+1}^\pi} \mathbb{E}_Q \left[\sum_{t=1}^T C_t^\pi \right]. \quad (6.33)$$

For all $s \in [0, T-1]$, $\pi \in \Pi_{s+1}$, and $Q \in \mathbf{MAR}_{s+1}^\pi$,

$$\mathbb{E}_Q \left[\sum_{t=T-s}^T C_t^\pi \right] = \sum_{\mathbf{q} \in \text{supp}(Q_{[T-s-1]})} \hat{V}^{s+1}(y_{T-s}^\pi(\mathbf{q}), q_{T-s-1}) Q_{[T-s-1]}(\mathbf{q}). \quad (6.34)$$

For all $s \in [0, T-1]$,

$$\inf_{\pi \in \Pi} \sup_{Q \in \mathbf{MAR}} \mathbb{E}_Q \left[\sum_{t=1}^T C_t^\pi \right] = \inf_{\pi \in \Pi_{s+1}} \sup_{Q \in \mathbf{MAR}} \mathbb{E}_Q \left[\sum_{t=1}^T C_t^\pi \right]. \quad (6.35)$$

Proof: We begin with the base case $s = 0$. (6.31) follows immediately from the fact that $\hat{f}^1(x, d) = C(x, d)$. (6.32) then follows from the fact that for all $\pi \in \Pi_0$, $Q \in \mathbf{MAR}_1^\pi$, and $\mathbf{q} \in \text{supp}(Q_{[T-1]})$, $Q_{T|\mathbf{q}} = \bar{Q}_{x_T^\pi(\mathbf{q}), q_{T-1}}^1 \in \hat{Q}^1(x_T^\pi(\mathbf{q}), q_{T-1})$, combined with the definition of \hat{g}^1 and \hat{Q}^1 .

We now prove (6.33). Let us fix any $\pi \in \Pi_0$. Note that for any $\epsilon > 0$, there exists $\tilde{Q}_\epsilon \in \mathbf{MAR}$ s.t. $\mathbb{E}[\sum_{t=1}^T C_t^\pi(D_{[t]}^{\tilde{Q}_\epsilon})] > \sup_{Q \in \mathbf{MAR}} \mathbb{E}_Q[\sum_{t=1}^T C_t^\pi] - \epsilon$. Let us fix any such $\epsilon > 0$ and corresponding $\tilde{Q}_\epsilon \in \mathbf{MAR}$, where we denote \tilde{Q}_ϵ by \tilde{Q} for clarity of exposition. We now prove that

$$\mathbb{E}\left[\sum_{t=1}^T C_t^\pi(D_{[t]}^{\tilde{Q}^{\pi,1}})\right] \geq \mathbb{E}\left[\sum_{t=1}^T C_t^\pi(D_{[t]}^{\tilde{Q}})\right]. \quad (6.36)$$

As $\tilde{Q}_{[T-1]}^{\pi,1} = \tilde{Q}_{[T-1]}$, it suffices by (6.31) to demonstrate that for every $\mathbf{q} \in \text{supp}(\tilde{Q}_{[T-1]})$,

$$\sum_{z \in [0, U]_{\mathcal{Q}}} \hat{f}^1(x_T^\pi(\mathbf{q}), z) \tilde{Q}_{T|\mathbf{q}}^{\pi,1}(z) \geq \sum_{z \in [0, U]_{\mathcal{Q}}} \hat{f}^1(x_T^\pi(\mathbf{q}), z) \tilde{Q}_{T|\mathbf{q}}(z). \quad (6.37)$$

By construction, $\tilde{Q}^{\pi,1} \in \mathbf{MAR}_1^\pi$, and thus for all $\mathbf{q} \in \text{supp}(\tilde{Q}_{[T-1]})$, $\tilde{Q}_{T|\mathbf{q}}^{\pi,1} = \bar{Q}_{x_T^\pi(\mathbf{q}), q_{T-1}}^1 \in \hat{Q}^1(x_T^\pi(\mathbf{q}), q_{T-1})$. Thus by definition, the left-hand side of (6.37) equals

$$\sup_{Q \in \mathfrak{M}(q_{T-1})} \sum_{z \in [0, U]_{\mathcal{Q}}} \hat{f}^1(x_T^\pi(\mathbf{q}), z) Q(z) = \hat{g}^1(x_T^\pi(\mathbf{q}), q_{T-1}). \quad (6.38)$$

Noting that the martingale property ensures $\tilde{Q}_{T|\mathbf{q}} \in \mathfrak{M}(q_{T-1})$ then completes the proof of (6.37), and (6.36). Letting $\epsilon \downarrow 0$ completes the proof of (6.33).

We now prove (6.34). As $\Pi_1 \subseteq \Pi_0$, it follows from (6.32), combined with the fact that $\pi \in \Pi_1$ implies $x_T^\pi(\mathbf{q}) = \bar{\Phi}_{y_T^\pi(\mathbf{q}), q_{T-1}}^1$, that

$$\mathbb{E}_Q[C_T^\pi] = \sum_{\mathbf{q} \in \text{supp}(Q_{[T-1]})} \hat{g}^1(\bar{\Phi}_{y_T^\pi(\mathbf{q}), q_{T-1}}^1, q_{T-1}) Q_{[T-1]}(\mathbf{q}).$$

As $\bar{\Phi}_{y_T^\pi(\mathbf{q}), q_{T-1}}^1 \in \hat{\Phi}^1(y_T^\pi(\mathbf{q}), q_{T-1})$, (6.34) then follows from definitions.

We now prove (6.35). For any $\epsilon > 0$, there exists $\tilde{\pi}_\epsilon \in \Pi_0$ s.t. $\inf_{\pi \in \Pi} \sup_{Q \in \mathbf{MAR}} \mathbb{E}_Q[\sum_{t=1}^T C_t^\pi] > \sup_{Q \in \mathbf{MAR}} \mathbb{E}_Q[\sum_{t=1}^T C_t^{\tilde{\pi}_\epsilon}] - \epsilon$. Let us fix any such $\epsilon > 0$ and corresponding $\tilde{\pi}_\epsilon \in \Pi_0$, where we denote $\tilde{\pi}_\epsilon$ by $\tilde{\pi}$ for clarity of exposition. We now prove that

$$\sup_{Q \in \mathbf{MAR}} \mathbb{E}_Q\left[\sum_{t=1}^T C_t^{\tilde{\pi}}\right] \geq \sup_{Q \in \mathbf{MAR}} \mathbb{E}_Q\left[\sum_{t=1}^T C_t^{\tilde{\pi}^1}\right]. \quad (6.39)$$

As $\Pi_1 \subseteq \Pi_0$, by (6.33) it suffices to demonstrate that

$$\sup_{Q \in \mathbf{MAR}_1^{\tilde{\pi}}} \mathbb{E}_Q\left[\sum_{t=1}^T C_t^{\tilde{\pi}}\right] \geq \sup_{Q \in \mathbf{MAR}_1^{\tilde{\pi}^1}} \mathbb{E}_Q\left[\sum_{t=1}^T C_t^{\tilde{\pi}^1}\right]. \quad (6.40)$$

Note that $Q \in \mathbf{MAR}_1^{\tilde{\pi}}$, combined with (6.32), implies that

$$\mathbb{E}_Q\left[\sum_{t=1}^T C_t^{\tilde{\pi}}\right] = \mathbb{E}_Q\left[\sum_{t=1}^{T-1} C_t^{\tilde{\pi}}\right] + \sum_{\mathbf{q} \in \text{supp}(Q_{[T-1]})} \hat{g}^1(x_T^{\tilde{\pi}}(\mathbf{q}), q_{T-1}) Q_{[T-1]}(\mathbf{q}).$$

$\tilde{\pi} \in \Pi_0$ implies that for all $Q \in \mathbf{MAR}_1^{\tilde{\pi}}$ and $\mathbf{q} \in [0, U]_{\mathcal{Q}}^{T-1}$, $x_T^{\tilde{\pi}}(\mathbf{q}) \geq y_T^{\tilde{\pi}}(\mathbf{q})$ and $x_T^{\tilde{\pi}}(\mathbf{q}) \in [0, U]_{\mathcal{Q}}$. By combining the above, we conclude from definitions that $Q \in \mathbf{MAR}_1^{\tilde{\pi}}$ implies

$$\begin{aligned} \mathbb{E}_Q \left[\sum_{t=1}^T C_t^{\tilde{\pi}} \right] &\geq \mathbb{E}_Q \left[\sum_{t=1}^{T-1} C_t^{\tilde{\pi}} + \sum_{\mathbf{q} \in \text{supp}(Q_{[T-1]})} \hat{V}^1(y_T^{\tilde{\pi}}(\mathbf{q}), q_{T-1}) Q_{[T-1]}(\mathbf{q}) \right] \\ &= \sum_{\mathbf{q} \in \text{supp}(Q_{[T-1]})} \left(\sum_{t=1}^{T-1} C(x_t^{\tilde{\pi}}(\mathbf{q}_{[t-1]}), q_t) + \hat{V}^1(y_T^{\tilde{\pi}}(\mathbf{q}), q_{T-1}) \right) Q_{[T-1]}(\mathbf{q}). \end{aligned} \quad (6.41)$$

Alternatively, if $Q \in \mathbf{MAR}_1^{\tilde{\pi}^1}$, it follows from (6.34), and the fact that $x_{[T-1]}^{\tilde{\pi}} = x_{[T-1]}^{\tilde{\pi}^1}$ and thus $y_T^{\tilde{\pi}^1} = y_T^{\tilde{\pi}}$, that

$$\mathbb{E}_Q \left[\sum_{t=1}^T C_t^{\tilde{\pi}^1} \right] = \sum_{\mathbf{q} \in \text{supp}(Q_{[T-1]})} \left(\sum_{t=1}^{T-1} C(x_t^{\tilde{\pi}^1}(\mathbf{q}_{[t-1]}), q_t) + \hat{V}^1(y_T^{\tilde{\pi}^1}(\mathbf{q}), q_{T-1}) \right) Q_{[T-1]}(\mathbf{q}). \quad (6.42)$$

For any $\epsilon_2 > 0$, there exists $\tilde{Q}_{\epsilon_2} \in \mathbf{MAR}_1^{\tilde{\pi}^1}$ (which we denote simply as \tilde{Q}) s.t.

$$\mathbb{E}_{\tilde{Q}} \left[\sum_{t=1}^T C_t^{\tilde{\pi}^1} \right] > \sup_{Q \in \mathbf{MAR}_1^{\tilde{\pi}^1}} \mathbb{E}_Q \left[\sum_{t=1}^T C_t^{\tilde{\pi}^1} \right] - \epsilon_2. \quad (6.43)$$

Lemma 17 implies the existence of $\tilde{Q}^{\tilde{\pi},1} \in \mathbf{MAR}_1^{\tilde{\pi}}$ s.t. $\tilde{Q}_{[T-1]}^{\tilde{\pi},1} = \tilde{Q}_{[T-1]}$. Trivially, $\sup_{Q \in \mathbf{MAR}_1^{\tilde{\pi}}} \mathbb{E}_Q \left[\sum_{t=1}^T C_t^{\tilde{\pi}} \right] \geq \mathbb{E}_{\tilde{Q}^{\tilde{\pi},1}} \left[\sum_{t=1}^T C_t^{\tilde{\pi}} \right]$, and it follows from (6.41), (6.42), and (6.43) that

$$\begin{aligned} \mathbb{E}_{\tilde{Q}^{\tilde{\pi},1}} \left[\sum_{t=1}^T C_t^{\tilde{\pi}} \right] &\geq \sum_{\mathbf{q} \in \text{supp}(\tilde{Q}_{[T-1]})} \left(\sum_{t=1}^{T-1} C(x_t^{\tilde{\pi}}(\mathbf{q}_{[t-1]}), q_t) + \hat{V}^1(y_T^{\tilde{\pi}}(\mathbf{q}), q_{T-1}) \right) \tilde{Q}_{[T-1]}(\mathbf{q}) \\ &= \mathbb{E}_{\tilde{Q}} \left[\sum_{t=1}^T C_t^{\tilde{\pi}^1} \right] > \sup_{Q \in \mathbf{MAR}_1^{\tilde{\pi}^1}} \mathbb{E}_Q \left[\sum_{t=1}^T C_t^{\tilde{\pi}^1} \right] - \epsilon_2. \end{aligned}$$

Letting $\epsilon_2 \downarrow 0$ completes the proof of (6.40) and (6.39). Letting $\epsilon \downarrow 0$ completes the proof of (6.35).

We now proceed by induction. Suppose that (6.31) - (6.35) hold for all $s' \in [0, s]$ for some $s \leq T-2$. We now prove that (6.31) - (6.35) also hold for $s+1$, and begin by verifying (6.31). Note that for all $\pi \in \Pi_{s+1}$ and $Q \in \mathbf{MAR}_{s+1}^{\pi}$,

$$\mathbb{E}_Q [C_{T-s-1}^{\pi}] = \sum_{\mathbf{q} \in \text{supp}(Q_{[T-s-2]})} \left(\sum_{z \in [0, U]_{\mathcal{Q}}} C(x_{T-s-1}^{\pi}(\mathbf{q}), z) Q_{T-s-1|\mathbf{q}}(z) \right) Q_{[T-s-2]}(\mathbf{q}). \quad (6.44)$$

Combining with (6.34), applied to s , and the fact that $y_{T-s}^{\pi}(\mathbf{q}; z) = x_{T-s-1}^{\pi}(\mathbf{q}) - z$, we conclude

$$\mathbb{E}_Q \left[\sum_{t=T-s}^T C_t^{\pi} \right] = \sum_{\mathbf{q} \in \text{supp}(Q_{[T-s-2]})} \left(\sum_{z \in [0, U]_{\mathcal{Q}}} \hat{V}^{s+1}(x_{T-s-1}^{\pi}(\mathbf{q}) - z, z) Q_{T-s-1|\mathbf{q}}(z) \right) Q_{[T-s-2]}(\mathbf{q}). \quad (6.45)$$

Combining (6.44) - (6.45) with the definition of \hat{f}^{s+2} completes the proof of (6.31).

We now prove (6.32). It follows from (6.31), applied to $s+1$, and the fact that $\mathbf{MAR}_{s+2}^{\pi} \subseteq \mathbf{MAR}_{s+1}^{\pi}$, that for all $\pi \in \Pi_{s+1}$ and $Q \in \mathbf{MAR}_{s+2}^{\pi}$,

$$\mathbb{E}_Q \left[\sum_{t=T-s-1}^T C_t^{\pi} \right] = \sum_{\mathbf{q} \in \text{supp}(Q_{[T-s-2]})} \left(\sum_{z \in [0, U]_{\mathcal{Q}}} \hat{f}^{s+2}(x_{T-s-1}^{\pi}(\mathbf{q}), z) Q_{T-s-1|\mathbf{q}}(z) \right) Q_{[T-s-2]}(\mathbf{q}).$$

As $Q \in \mathbf{MAR}_{s+2}^\pi$, $\mathbf{q} \in \text{supp}(Q_{[T-s-2]})$ implies $Q_{T-s-1|\mathbf{q}} = \overline{Q}_{x_{T-s-1}^\pi(\mathbf{q}), q_{T-s-2}}^{s+2} \in \hat{Q}^{s+2}(x_{T-s-1}^\pi(\mathbf{q}), q_{T-s-2})$, (6.32) then follows from the definitions of \hat{g}^{s+2} and \hat{Q}^{s+2} .

We now prove (6.33). As $\Pi_{s+1} \subseteq \Pi_s$, it follows from (6.33), applied to s , that it suffices to demonstrate that for all $\pi \in \Pi_{s+1}$,

$$\sup_{Q \in \mathbf{MAR}_{s+1}^\pi} \mathbb{E}_Q \left[\sum_{t=1}^T C_t^\pi \right] = \sup_{Q \in \mathbf{MAR}_{s+2}^\pi} \mathbb{E}_Q \left[\sum_{t=1}^T C_t^\pi \right]. \quad (6.46)$$

Let us fix any $\pi \in \Pi_{s+1}$. For any $\epsilon > 0$, there exists $\tilde{Q}_\epsilon \in \mathbf{MAR}_{s+1}^\pi$ s.t. $\mathbb{E}[\sum_{t=1}^T C_t^\pi(D_{[t]}^{\tilde{Q}_\epsilon})] > \sup_{Q \in \mathbf{MAR}_{s+1}^\pi} \mathbb{E}_Q[\sum_{t=1}^T C_t^\pi] - \epsilon$. Let us fix any such $\epsilon > 0$ and corresponding $\tilde{Q}_\epsilon \in \mathbf{MAR}_{s+1}^\pi$, where we denote \tilde{Q}_ϵ by \tilde{Q} for clarity of exposition. We now prove that

$$\mathbb{E} \left[\sum_{t=1}^T C_t^\pi(D_{[t]}^{\tilde{Q}^{\pi, s+2}}) \right] \geq \mathbb{E} \left[\sum_{t=1}^T C_t^\pi(D_{[t]}^{\tilde{Q}}) \right]. \quad (6.47)$$

As $\tilde{Q}_{[T-s-2]}^{\pi, s+2} = \tilde{Q}_{[T-s-2]}$, by (6.31), applied to $s+1$, it suffices to demonstrate that for every $\mathbf{q} \in \text{supp}(\tilde{Q}_{[T-s-2]})$,

$$\sum_{z \in [0, U]_{\mathcal{Q}}} \hat{f}^{s+2}(x_{T-s-1}^\pi(\mathbf{q}), z) \tilde{Q}_{T-s-1|\mathbf{q}}^{\pi, s+2}(z) \geq \sum_{z \in [0, U]_{\mathcal{Q}}} \hat{f}^{s+2}(x_{T-s-1}^\pi(\mathbf{q}), z) \tilde{Q}_{T-s-1|\mathbf{q}}(z). \quad (6.48)$$

By construction, $\tilde{Q}^{\pi, s+2} \in \mathbf{MAR}_{s+2}^\pi$, and thus for all $\mathbf{q} \in \text{supp}(\tilde{Q}_{[T-s-2]})$, $\tilde{Q}_{T-s-1|\mathbf{q}}^{\pi, s+2} = \overline{Q}_{x_{T-s-1}^\pi(\mathbf{q}), q_{T-s-2}}^{s+2} \in \hat{Q}^{s+2}(x_{T-s-1}^\pi(\mathbf{q}), q_{T-s-2})$. Thus by definition, the left-hand side of (6.48) equals

$$\sup_{Q \in \mathfrak{M}(q_{T-s-2})} \sum_{z \in [0, U]_{\mathcal{Q}}} \hat{f}^{s+2}(x_{T-s-1}^\pi(\mathbf{q}), z) Q(z) = \hat{g}^{s+2}(x_{T-s-1}^\pi(\mathbf{q}), q_{T-s-2}). \quad (6.49)$$

Noting that the martingale property ensures $\tilde{Q}_{T-s-1|\mathbf{q}} \in \mathfrak{M}(q_{T-s-2})$ then completes the proof of (6.48), and (6.47). Letting $\epsilon \downarrow 0$ further completes the proof of (6.33).

We now prove (6.34). As $\Pi_{s+2} \subseteq \Pi_{s+1}$, it follows from (6.32), applied to $s+1$, and the fact that $\pi \in \Pi_{s+2}$ implies $x_{T-s-1}^\pi(\mathbf{q}) = \overline{\Phi}_{y_{T-s-1}^\pi(\mathbf{q}), q_{T-s-2}}^{s+2}$, that

$$\mathbb{E}_Q \left[\sum_{t=T-s-1}^T C_t^\pi \right] = \sum_{\mathbf{q} \in \text{supp}(Q_{[T-s-2]})} \hat{g}^{s+2}(\overline{\Phi}_{y_{T-s-1}^\pi(\mathbf{q}), q_{T-s-2}}^{s+2}) Q_{[T-s-2]}(\mathbf{q}).$$

As $\overline{\Phi}_{y_{T-s-1}^\pi(\mathbf{q}), q_{T-s-2}}^{s+2} \in \hat{\Phi}^{s+2}(y_{T-s-1}^\pi(\mathbf{q}), q_{T-s-2})$, (6.34) then follows from definitions.

We now prove (6.35). First, observe that by (6.35), applied to s , it suffices to prove that

$$\inf_{\pi \in \Pi_{s+1}} \sup_{Q \in \mathbf{MAR}} \mathbb{E}_Q \left[\sum_{t=1}^T C_t^\pi \right] = \inf_{\pi \in \Pi_{s+2}} \sup_{Q \in \mathbf{MAR}} \mathbb{E}_Q \left[\sum_{t=1}^T C_t^\pi \right]. \quad (6.50)$$

For any $\epsilon > 0$, there exists $\tilde{\pi}_\epsilon \in \Pi_{s+1}$ s.t. $\inf_{\pi \in \Pi_{s+1}} \sup_{Q \in \mathbf{MAR}} \mathbb{E}_Q[\sum_{t=1}^T C_t^\pi] > \sup_{Q \in \mathbf{MAR}} \mathbb{E}_Q[\sum_{t=1}^T C_t^{\tilde{\pi}_\epsilon}] - \epsilon$. Let us fix any such $\epsilon > 0$ and corresponding $\tilde{\pi}_\epsilon \in \Pi_{s+1}$, where we denote $\tilde{\pi}_\epsilon$ by $\tilde{\pi}$ for clarity of exposition. We now prove that

$$\sup_{Q \in \mathbf{MAR}} \mathbb{E}_Q \left[\sum_{t=1}^T C_t^{\tilde{\pi}} \right] \geq \sup_{Q \in \mathbf{MAR}} \mathbb{E}_Q \left[\sum_{t=1}^T C_t^{\tilde{\pi}^{s+2}} \right].$$

As $\Pi_{s+2} \subseteq \Pi_{s+1}$, by (6.33), applied to $s+1$, it suffices to demonstrate that

$$\sup_{Q \in \mathbf{MAR}_{s+2}^{\tilde{\pi}}} \mathbb{E}_Q \left[\sum_{t=1}^T C_t^{\tilde{\pi}} \right] \geq \sup_{Q \in \mathbf{MAR}_{s+2}^{\tilde{\pi}^{s+2}}} \mathbb{E}_Q \left[\sum_{t=1}^T C_t^{\tilde{\pi}^{s+2}} \right]. \quad (6.51)$$

Note that $Q \in \mathbf{MAR}_{s+2}^{\tilde{\pi}}$, combined with (6.32), applied to $s+1$, implies that

$$\mathbb{E}_Q \left[\sum_{t=1}^T C_t^{\tilde{\pi}} \right] = \mathbb{E}_Q \left[\sum_{t=1}^{T-s-2} C_t^{\tilde{\pi}} \right] + \sum_{\mathbf{q} \in \text{supp}(Q_{[T-s-2]})} \hat{g}^{s+2}(x_{T-s-1}^{\tilde{\pi}}(\mathbf{q}), q_{T-s-2}) Q_{[T-s-2]}(\mathbf{q}).$$

$\tilde{\pi} \in \Pi_{s+1}$ implies that for all $Q \in \mathbf{MAR}_{s+2}^{\tilde{\pi}}$ and $\mathbf{q} \in [0, U]_{\mathcal{Q}}^{T-s-2}$, $x_{T-s-1}^{\tilde{\pi}}(\mathbf{q}) \geq y_{T-s-1}^{\tilde{\pi}}(\mathbf{q})$ and $x_{T-s-1}^{\tilde{\pi}}(\mathbf{q}) \in [0, U]_{\mathcal{Q}}$. By combining the above, we conclude from definitions that $Q \in \mathbf{MAR}_{s+2}^{\tilde{\pi}}$ implies $\mathbb{E}_Q \left[\sum_{t=1}^T C_t^{\tilde{\pi}} \right]$ is at least

$$\mathbb{E}_Q \left[\sum_{t=1}^{T-s-2} C_t^{\tilde{\pi}} \right] + \sum_{\mathbf{q} \in \text{supp}(Q_{[T-s-2]})} \hat{V}^{s+2}(y_{T-s-1}^{\tilde{\pi}}(\mathbf{q}), q_{T-s-2}) Q_{[T-s-2]}(\mathbf{q}),$$

which is itself equal to

$$\sum_{\mathbf{q} \in \text{supp}(Q_{[T-s-2]})} \left(\sum_{t=1}^{T-s-2} C(x_t^{\tilde{\pi}}(\mathbf{q}_{[t-1]}), q_t) + \hat{V}^{s+2}(y_{T-s-1}^{\tilde{\pi}}(\mathbf{q}), q_{T-s-2}) \right) Q_{[T-s-2]}(\mathbf{q}). \quad (6.52)$$

Alternatively, if $Q \in \mathbf{MAR}_{s+2}^{\tilde{\pi}^{s+2}}$, it follows from (6.34), applied to $s+1$, and the fact that $x_{[T-s-2]}^{\tilde{\pi}} = x_{[T-s-2]}^{\tilde{\pi}^{s+2}}$, that $\mathbb{E}_Q \left[\sum_{t=1}^T C_t^{\tilde{\pi}^{s+2}} \right]$ equals

$$\mathbb{E}_Q \left[\sum_{t=1}^{T-s-2} C_t^{\tilde{\pi}} \right] + \sum_{\mathbf{q} \in \text{supp}(Q_{[T-s-2]})} \hat{V}^{s+2}(y_{T-s-1}^{\tilde{\pi}}(\mathbf{q}), q_{T-s-2}) Q_{[T-s-2]}(\mathbf{q}),$$

which is itself equal to

$$\sum_{\mathbf{q} \in \text{supp}(Q_{[T-s-2]})} \left(\sum_{t=1}^{T-s-2} C(x_t^{\tilde{\pi}}(\mathbf{q}_{[t-1]}), q_t) + \hat{V}^{s+2}(y_{T-s-1}^{\tilde{\pi}}(\mathbf{q}), q_{T-s-2}) \right) Q_{[T-s-2]}(\mathbf{q}). \quad (6.53)$$

For any $\epsilon_2 > 0$, there exists $\tilde{Q}_{\epsilon_2} \in \mathbf{MAR}_{s+2}^{\tilde{\pi}^{s+2}}$ (which we denote simply as \tilde{Q}) s.t.

$$\mathbb{E}_{\tilde{Q}} \left[\sum_{t=1}^T C_t^{\tilde{\pi}^{s+2}} \right] > \sup_{Q \in \mathbf{MAR}_{s+2}^{\tilde{\pi}^{s+2}}} \mathbb{E}_Q \left[\sum_{t=1}^T C_t^{\tilde{\pi}^{s+2}} \right] - \epsilon_2. \quad (6.54)$$

Lemma 17 implies the existence of $\tilde{Q}^{\tilde{\pi}, s+2} \in \mathbf{MAR}_{s+2}^{\tilde{\pi}}$ s.t. $\tilde{Q}_{[T-s-2]}^{\tilde{\pi}, s+2} = \tilde{Q}_{[T-s-2]}$. Trivially, $\sup_{Q \in \mathbf{MAR}_{s+2}^{\tilde{\pi}}} \mathbb{E}_Q \left[\sum_{t=1}^T C_t^{\tilde{\pi}} \right] \geq \mathbb{E}_{\tilde{Q}^{\tilde{\pi}, s+2}} \left[\sum_{t=1}^T C_t^{\tilde{\pi}} \right]$, which by (6.52), (6.53), and (6.54) is itself at least

$$\begin{aligned} & \sum_{\mathbf{q} \in \text{supp}(\tilde{Q}_{[T-s-2]}^{\tilde{\pi}, s+2})} \left(\sum_{t=1}^{T-s-2} C(x_t^{\tilde{\pi}}(\mathbf{q}_{[t-1]}), q_t) + \hat{V}^{s+2}(y_{T-s-1}^{\tilde{\pi}}(\mathbf{q}), q_{T-s-2}) \right) \tilde{Q}_{[T-s-2]}^{\tilde{\pi}, s+2}(\mathbf{q}) \\ &= \sum_{\mathbf{q} \in \text{supp}(\tilde{Q}_{[T-s-2]})} \left(\sum_{t=1}^{T-s-2} C(x_t^{\tilde{\pi}}(\mathbf{q}_{[t-1]}), q_t) + \hat{V}^{s+2}(y_{T-s-1}^{\tilde{\pi}}(\mathbf{q}), q_{T-s-2}) \right) \tilde{Q}_{[T-s-2]}(\mathbf{q}) \\ &= \mathbb{E}_{\tilde{Q}} \left[\sum_{t=1}^T C_t^{\tilde{\pi}^{s+2}} \right] > \sup_{Q \in \mathbf{MAR}_{s+2}^{\tilde{\pi}^{s+2}}} \mathbb{E}_Q \left[\sum_{t=1}^T C_t^{\tilde{\pi}^{s+2}} \right] - \epsilon_2. \end{aligned}$$

Letting $\epsilon_2 \downarrow 0$ completes the proof of (6.51) and (6.50). Letting $\epsilon \downarrow 0$ completes the proof of (6.35).

Combining all of the above completes the desired induction and proof of the theorem. \square .

6.3. Proof of Lemma 4

Proof: [Proof of Lemma 4] We first prove continuity. That $\mathbf{g}^T(x, d)$ is a continuous function of x on $(0, U) \setminus \bigcup_{i=\Gamma_d^{T-1}}^{T-1} \{B_i^T\}$, and a right-continuous function of x on $[0, U] \setminus \{U\}$, follows from definitions, and the fact that $F_i^T(x, d)$ and $G_i^T(x, d)$ are continuous functions of x on $[0, U]$ for all i . It similarly follows that $\lim_{x \uparrow B_i^T} \mathbf{g}^T(x, d)$ exists for all $i \in [\Gamma_d^{T-1}, T] \setminus \{0\}$. It thus suffices to demonstrate that $\lim_{x \uparrow B_i^T} \mathbf{g}^T(x, d)$ equals $\mathbf{g}^T(B_i^T, d)$ for all $i \in [\Gamma_d^{T-1}, T] \setminus \{0\}$. We treat two cases: $i = \Gamma_d^{T-1}$, and $i \in [\Gamma_d^{T-1} + 1, T]$, and begin with the case $i = \Gamma_d^{T-1}$. By assumption we preclude the case $i = 0$. Thus suppose $i = \Gamma_d^{T-1} \in [1, T]$. In this case,

$$\begin{aligned} \lim_{x \uparrow B_i^T} \mathbf{g}^T(x, d) &= F_{i-1}^T(B_i^T, d) \\ &= -bB_i^T + (b+T)B_i^T + (Tb - (b+1)i)d. \\ &= TB_i^T + (Tb - (b+1)i)d. \end{aligned}$$

Alternatively,

$$\begin{aligned} \mathbf{g}^T(B_i^T, d) &= G_i^T(B_i^T, d) \\ &= (T - \frac{b+T}{A_i^T}d)B_i^T + (T-i)bd \\ &= TB_i^T + (Tb - (b+1)i)d. \end{aligned}$$

Combining the above completes the proof for this case. Next, suppose $i \in [\Gamma_d^{T-1} + 1, T]$. In this case,

$$\begin{aligned} \lim_{x \uparrow B_i^T} \mathbf{g}^T(x, d) &= G_{i-1}^T(B_i^T, d) \\ &= (T - \frac{b+T}{A_{i-1}^T}d)B_i^T + (T - (i-1))bd \\ &= TB_i^T + (Tb - (b+1)i)d. \end{aligned}$$

Alternatively,

$$\begin{aligned} \mathbf{g}^T(B_i^T, d) &= G_i^T(B_i^T, d) \\ &= (T - \frac{b+T}{A_i^T}d)B_i^T + (T-i)bd = TB_i^T + (Tb - (b+1)i)d, \end{aligned}$$

completing the proof.

We now prove convexity. As in our proof of continuity, that $\mathbf{g}^T(x, d)$ is a right-differentiable function of x on $(0, U) \setminus \bigcup_{i=\Gamma_d^{T-1}}^{T-1} \{B_i^T\}$, with non-decreasing right-derivative on the same set, follows from definitions and the fact that $F_i^T(x, d)$ and $G_i^T(x, d)$ are linear functions of x on $[0, U]$ for all i . It similarly follows that $\mathbf{g}^T(x, d)$ is a right-differentiable function of x on $[0, U] \setminus \{U\}$, and that $\lim_{x \uparrow B_i^T} \partial_x^+ \mathbf{g}^T(x, d)$ exists for all $i \in [\Gamma_d^{T-1}, T-1] \setminus \{0\}$. It thus suffices to demonstrate that $\lim_{x \uparrow B_i^T} \partial_x^+ \mathbf{g}^T(x, d) \leq \partial_x^+ \mathbf{g}^T(B_i^T, d)$ for all $i \in [\Gamma_d^{T-1}, T-1] \setminus \{0\}$. We treat two cases: $i = \Gamma_d^{T-1}$, and $i \in [\Gamma_d^{T-1} + 1, T-1]$, and begin with the case $i = \Gamma_d^{T-1}$. By assumption we preclude the cases $i = 0, T$. Thus suppose $i = \Gamma_d^{T-1} \in [1, T-1]$. Then

$$\begin{aligned} \lim_{x \uparrow B_i^T} \partial_x^+ \mathbf{g}^T(x, d) &= \lim_{x \uparrow B_i^T} \partial_x^+ F_{i-1}^T(x, d) \\ &= -b. \end{aligned}$$

Alternatively,

$$\begin{aligned}\partial_x^+ \mathbf{g}^T(B_i^T, d) &= \partial_x^+ G_i^T(B_i^T, d) \\ &= T - \frac{b+T}{A_i^T} d. \\ &\geq T - \frac{b+T}{A_i^T} A_i^T = -b.\end{aligned}$$

Combining the above completes the proof for this case. Next, suppose $i \in [\Gamma_d^{T-1} + 1, T - 1]$. Then

$$\begin{aligned}\lim_{x \uparrow B_i^T} \partial_x^+ \mathbf{g}^T(x, d) &= \lim_{x \uparrow B_i^T} \partial_x^+ G_{i-1}^T(x, d) \\ &= T - \frac{b+T}{A_{i-1}^T} d.\end{aligned}$$

Alternatively,

$$\begin{aligned}\partial_x^+ \mathbf{g}^T(B_i^T, d) &= \partial_x^+ G_i^T(B_i^T, d) \\ &= T - \frac{b+T}{A_i^T} d.\end{aligned}$$

The desired result then follows from the fact that A_i^T is increasing in i and d is non-negative. Combining the above completes the proof. \square

6.4. Proofs of Lemmas 5 and 6

Proof: [Proof of Lemma 5] The desired result follows directly from the observation that A_{j+1}^{T+1} is trivially strictly less than A_{j+1}^T for all $j \in [0, T - 2]$, and $\frac{A_j^T}{A_{j+1}^{T+1}} = \frac{\frac{j+1}{b+j+1}}{\frac{T}{b+T}} < 1$ for all $j \in [0, T - 2]$. \square

Proof: [Proof of Lemma 6] Let $i = \Gamma_d^T, j = \Gamma_d^{T-1}$. From Lemma 4, it suffices to prove that $\partial_x^+ \mathbf{g}^T(x, d) \leq 0$ for all $x < B_i^T$, and $\partial_x^+ \mathbf{g}^T(B_i^T, d) \geq 0$; or that $\partial_x^+ \mathbf{g}^T(x, d) \leq 0$ for all $x < U$ and $B_i^T = U$. It follows from Lemma 5 that $i \in \{j, j + 1\}$. We now proceed by a case analysis. First, suppose $i = j$. In that case, $B_i^T = B_j^T, j \leq T - 1$, and $B_i^T < U$. We conclude that for all $x < B_i^T$,

$$\partial_x^+ \mathbf{g}^T(x, d) = \partial_x^+ F_{j-1}^T(x, \mu) = -b.$$

Noting that

$$\begin{aligned}\partial_x^+ \mathbf{g}^T(B_i^T, d) &= \partial_x^+ G_i^T(B_i^T, d) \\ &= T - \frac{b+T}{A_i^T} d \geq T - \frac{b+T}{A_i^T} A_i^{T+1} = 0\end{aligned}$$

completes the proof in this setting.

Alternatively, suppose that $i = j + 1$. In this case, $B_i^T > B_j^T$, and

$$\begin{aligned}\lim_{x \uparrow B_i^T} \partial_x^+ \mathbf{g}^T(x, d) &= \lim_{x \uparrow B_i^T} \partial_x^+ G_{i-1}^T(x, d) \\ &= T - \frac{b+T}{A_{i-1}^T} d \leq T - \frac{b+T}{A_{i-1}^T} A_{i-1}^{T+1} = 0.\end{aligned}$$

If $B_i^T = U$, the result follows from Lemma 4. Otherwise,

$$\begin{aligned}\partial_x^+ \mathbf{g}^T(B_i^T, d) &= \partial_x^+ G_i^T(B_i^T, d) \\ &= T - \frac{b+T}{A_i^T} d \geq T - \frac{b+T}{A_i^T} A_i^{T+1} = 0.\end{aligned}$$

Combining the above completes the proof. \square

6.5. Proof of Lemma 7

Proof: [Proof of Lemma 7] Let us proceed by showing that for each fixed $x \in [0, U]$, $\bar{\mathbf{g}}^T(x, d) = \mathbf{g}^T\left(\max(\beta_d^T, x - d), d\right)$ for all $d \in [0, U]$. As the equivalence is easily verified for the case $d = 0$, in which case $\bar{\mathbf{g}}^T(x, d) = \mathbf{g}^T\left(\max(\beta_d^T, x - d), d\right) = Tx$, suppose $d > 0$. We proceed by a case analysis, beginning with the setting $d \in (0, z_x^T)$. In this case, $\max(\beta_d^T, x - d) = x - d$, and

$$x > d + B_{\Gamma_d^T}^T \geq d + B_{\Gamma_{d-1}^T}^T,$$

where the second inequality follows from Lemma 5 and the monotonicity (in i) of B_i^T . Combining with the easily verified fact that $\bar{F}_j^T(x, d) = G_j^T(x - d, d)$ for all j completes the proof in this case.

Alternatively, suppose $d \in [z_x^T, U]$, which implies that $\max(\beta_d^T, x - d) = \beta_d^T$, and $d \geq x - \beta_d^T$. We proceed by a case analysis. First, suppose $d \in (A_j^T, A_{j+1}^T]$ for some $j \in [-1, T - 2]$. In this case, Lemma 5 implies that $d \in (A_j^T, A_{j+1}^T] \cap (A_j^{T+1}, A_{j+1}^{T+1}]$, and

$$\begin{aligned} \bar{\mathbf{g}}^T(x, d) &= \bar{G}_j^T(d). \\ &= G_{j+1}^T(B_{j+1}^T, d) = \mathbf{g}^T(\beta_d^T, d). \end{aligned}$$

Alternatively, suppose $d \in (A_{j+1}^T, A_{j+2}^T]$ for some $j \in [-1, T - 2]$. In this case, Lemma 5 implies that $d \in (A_j^T, A_{j+1}^T] \cap (A_{j+1}^{T+1}, A_{j+2}^{T+1}]$, and

$$\begin{aligned} \bar{\mathbf{g}}^T(x, d) &= \bar{G}_{j+1}^T(d). \\ &= G_{j+2}^T(B_{j+2}^T, d) = \mathbf{g}^T(\beta_d^T, d). \end{aligned}$$

Lemma 5 implies that this treats all cases. Combining the above completes the proof. \square

6.6. Proof of Lemma 8

Proof: [Proof of Lemma 8] First, let us treat the case $d \in [0, z_x^T)$, and begin by proving continuity. Right-continuity at 0 when $x \neq 0$ follows from the fact that $\lim_{d \downarrow 0} \bar{F}_j^T(x, d) = Tx$ for all j , and right-continuity at 0 when $x = 0$ follows from definitions. That $\bar{\mathbf{g}}^T(x, d)$ is a continuous function of d on $(0, z_x^T) \setminus \bigcup_{j=1}^{T-1} \{x - B_j^T\}$, and a left-continuous function of d on $(0, z_x^T)$, follows from the continuity (in d) of $\bar{F}_j^T(x, d)$ for all j . It similarly follows that $\lim_{d \downarrow x - B_j^T} \bar{\mathbf{g}}^T(x, d)$ exists for all j s.t. $x - B_j^T \in (0, z_x^T)$, and it thus suffices to demonstrate that $\lim_{d \downarrow x - B_j^T} \bar{\mathbf{g}}^T(x, d)$ equals $\bar{\mathbf{g}}^T(x, x - B_j^T)$ for all $j \in [1, T - 1]$ s.t. $x - B_j^T \in (0, z_x^T)$. Note that for all such j ,

$$\begin{aligned} \lim_{d \downarrow x - B_j^T} \bar{\mathbf{g}}^T(x, d) &= \bar{F}_{j-1}^T(x, x - B_j^T) \\ &= Tx + ((b-1)T - b(j-1) - \frac{b+T}{A_{j-1}^T}x)(x - B_j^T) + \frac{b+T}{A_{j-1}^T}(x - B_j^T)^2. \end{aligned}$$

Alternatively,

$$\begin{aligned} \bar{\mathbf{g}}^T(x, x - B_j^T) &= \bar{F}_j^T(x, x - B_j^T) \\ &= Tx + ((b-1)T - bj - \frac{b+T}{A_j^T}x)(x - B_j^T) + \frac{b+T}{A_j^T}(x - B_j^T)^2. \end{aligned}$$

It follows that $\bar{\mathbf{g}}^T(x, x - B_j^T) - \lim_{d \downarrow x - B_j^T} \bar{\mathbf{g}}^T(x, d)$ equals

$$\left(-b + \left(\frac{b+T}{A_{j-1}^T} - \frac{b+T}{A_j^T}\right)x\right)(x - B_j^T) - \left(\frac{b+T}{A_{j-1}^T} - \frac{b+T}{A_j^T}\right)(x - B_j^T)^2$$

$$\begin{aligned}
&= (x - B_j^T) \left(-b + \left(\frac{b+T}{A_{j-1}^T} - \frac{b+T}{A_j^T} \right) x - \left(\frac{b+T}{A_{j-1}^T} - \frac{b+T}{A_j^T} \right) (x - B_j^T) \right) \\
&= (x - B_j^T) \left(-b + (b+T) B_j^T \left(\frac{1}{A_{j-1}^T} - \frac{1}{A_j^T} \right) \right) = 0,
\end{aligned}$$

completing the proof of continuity.

We now prove convexity. Again applying Lemma 3, it suffices to demonstrate that $\partial_d^+ \bar{\mathbf{g}}^T(x, d)$ exists and is non-decreasing on $(0, z_x^T)$. Since $\bar{F}_j^T(x, d)$ is a convex quadratic function of d for all j , we conclude that: $\partial_d^+ \bar{\mathbf{g}}^T(x, d)$ exists on $(0, z_x^T)$; $\partial_d^+ \bar{\mathbf{g}}^T(x, d)$ is non-decreasing on $(0, z_x^T) \setminus \bigcup_{j=1}^{T-1} \{x - B_j^T\}$; and $\lim_{d \uparrow x - B_j^T} \partial_d^+ \bar{\mathbf{g}}^T(x, d)$ exists for all $j \in [1, T-1]$ s.t. $x - B_j^T \in (0, z_x^T)$. It thus suffices to demonstrate that $\lim_{d \uparrow x - B_j^T} \partial_d^+ \bar{\mathbf{g}}^T(x, d) \leq \partial_d^+ \bar{\mathbf{g}}^T(x, x - B_j^T)$ for all $j \in [1, T-1]$ s.t. $x - B_j^T \in (0, z_x^T)$. Note that for any such j ,

$$\begin{aligned}
\lim_{d \uparrow x - B_j^T} \partial_d^+ \bar{\mathbf{g}}^T(x, d) &= \partial_d^+ \bar{F}_j^T(x, x - B_j^T) \\
&= (b-1)T - bj - \frac{b+T}{A_j^T} x + 2 \frac{b+T}{A_j^T} (x - B_j^T) \\
&= (b-1)T - (b+2)j + x \frac{b+T}{A_j^T}.
\end{aligned}$$

Alternatively, it follows from continuity that

$$\begin{aligned}
\partial_d^+ \bar{\mathbf{g}}^T(x, x - B_j^T) &= \partial_d^+ \bar{F}_{j-1}^T(x, x - B_j^T) \\
&= (b-1)T - b(j-1) - \frac{b+T}{A_{j-1}^T} x + 2 \frac{b+T}{A_{j-1}^T} (x - B_j^T) \\
&= (b-1)T - (b+2)j - b + x \frac{b+T}{A_{j-1}^T}.
\end{aligned}$$

It follows that

$$\partial_d^+ \bar{\mathbf{g}}^T(x, x - B_j^T) - \lim_{d \uparrow x - B_j^T} \partial_d^+ \bar{\mathbf{g}}^T(x, d) \quad (6.55)$$

equals

$$-b + x \frac{b+T}{A_{j-1}^T} - x \frac{b+T}{A_j^T},$$

which will be the same sign as

$$-bA_{j-1}^T + x(b+T) - x(b+T) \frac{j}{b+j} = -bA_{j-1}^T + x(b+T) \frac{b}{b+j}. \quad (6.56)$$

Noting that $\frac{b+j}{b+T} A_{j-1}^T = B_j^T$, and multiplying through the right-hand side of (6.56) by $\frac{b+j}{b(b+T)}$, we further conclude that (6.55) will be the same sign as $x - B_j^T$. Since by assumption $x - B_j^T \in (0, z_x^T)$, we conclude that $x - B_j^T \geq 0$, completing the proof of continuity and convexity for $d \in (0, z_x^T)$ and right-continuity at 0.

Next, let us treat the case $d \in (z_x^T, U]$, and begin by proving continuity. That $\bar{\mathbf{g}}^T(x, d)$ is a continuous function of d on $(z_x^T, U) \setminus \bigcup_{j=0}^{T-1} \{A_j^{T+1}\}$, and a left-continuous function of d on $(z_x^T, U]$, follows from the continuity (in d) of $\bar{G}_j^T(d)$ for all j . It similarly follows that $\lim_{d \downarrow A_j^{T+1}} \bar{\mathbf{g}}^T(x, d)$ exist for all $j \in [0, T-1]$ s.t. $A_j^{T+1} > z_x^T$. It thus suffices to demonstrate that $\lim_{d \downarrow A_j^{T+1}} \bar{\mathbf{g}}^T(x, d) = \bar{\mathbf{g}}^T(x, A_j^{T+1})$ for all $j \in [0, T-1]$ s.t. $A_j^{T+1} > z_x^T$. Note that for any such j ,

$$\begin{aligned}
\lim_{d \downarrow A_j^{T+1}} \bar{\mathbf{g}}^T(x, d) &= \bar{G}_j^T(A_j^{T+1}) \\
&= TB_{j+1}^T + (Tb - (b+1)(j+1))A_j^{T+1}.
\end{aligned}$$

Alternatively,

$$\begin{aligned}\bar{\mathbf{g}}^T(x, A_j^{T+1}) &= \bar{G}_{j-1}^T(A_j^{T+1}) \\ &= TB_j^T + (Tb - (b+1)j)A_j^{T+1}.\end{aligned}$$

Thus

$$\begin{aligned}\bar{\mathbf{g}}^T(x, A_j^{T+1}) - \lim_{d \downarrow A_j^{T+1}} \bar{\mathbf{g}}^T(x, d) &= T(B_j^T - B_{j+1}^T) + (b+1)A_j^{T+1} \\ &= \frac{T}{b+T}(jA_j^T - (j+1)A_{j+1}^T) + (b+1)\frac{T}{b+T}A_j^T \\ &= \frac{T}{b+T}((b+j+1)A_j^T - (j+1)A_{j+1}^T) = 0,\end{aligned}$$

completing the proof of continuity.

We now prove concavity. Again applying Lemma 3, it suffices to demonstrate that $\partial_d^+ \bar{\mathbf{g}}^T(x, d)$ exists and is non-increasing on (z_x^T, U) . Since $\bar{G}_j^T(d)$ is a linear function of d for all j , it follows from the piece-wise definition of $\bar{\mathbf{g}}^T(x, d)$ that demonstrating the desired concavity is equivalent to showing that $\partial_d^+ \bar{G}_j^T(0)$ is non-increasing in j . Noting that $\partial_d^+ \bar{G}_j^T(0) = Tb - (b+1)(j+1)$, which is trivially decreasing in j , completes the proof.

Finally, let us prove continuity at z_x^T . We consider two cases, depending on how $\eta_d \triangleq x - d$ comes to go from lying above β_d^T to lying below β_d^T . This ‘‘crossing’’ can occur in two ways. In particular, either β_d^T and η_d actually intersect, or z_x^T occurs at a jump discontinuity of β_d^T and the two functions never truly intersect. We proceed by a case analysis. Let $i = \Gamma_{z_x^T}^T$.

First, suppose that β_d^T and η_d actually intersect at z_x^T , namely $B_i^T = x - z_x^T$. If $z_x^T \in \bigcup_{j=-1}^T \{A_j^{T+1}\}$, the proof of right-continuity follows identically to our previous proof of right-continuity at A_j^{T+1} for all $j \in [0, T-1]$ s.t. $A_j^{T+1} > z_x^T$, and we omit the details. Otherwise, right-continuity at z_x^T follows from definitions. Either way, we need only demonstrate left-continuity. Note that $z_x^T \in (x - B_{i+1}^T, x - B_i^T]$, and

$$\begin{aligned}\lim_{d \uparrow z_x^T} \bar{\mathbf{g}}^T(x, d) &= \bar{F}_i^T(x, z_x^T) \\ &= Tx + ((b-1)T - bi - \frac{b+T}{A_i^T}x)(x - B_i^T) + \frac{b+T}{A_i^T}(x - B_i^T)^2 \\ &= TB_i^T + (Tb - (b+1)i)(x - B_i^T).\end{aligned}$$

Alternatively,

$$\begin{aligned}\bar{\mathbf{g}}^T(x, z_x^T) &= \bar{G}_{i-1}^T(x - B_i^T) \\ &= TB_i^T + (Tb - (b+1)i)(x - B_i^T),\end{aligned}$$

completing the proof of continuity in this case.

Alternatively, suppose that β_d^T and η_d do not truly intersect at z_x^T . In this case, $z_x^T = A_i^{T+1} \in (x - B_{i+1}^T, x - B_i^T]$, and

$$\begin{aligned}\lim_{d \uparrow z_x^T} \bar{\mathbf{g}}^T(x, d) &= \bar{F}_i^T(x, A_i^{T+1}) \\ &= Tx + ((b-1)T - bi - \frac{b+T}{A_i^T}x)A_i^{T+1} + \frac{b+T}{A_i^T}(A_i^{T+1})^2 \\ &= b(T-i)A_i^{T+1},\end{aligned}$$

where the final equality follows from straightforward algebraic manipulations, the details of which we omit. Alternatively,

$$\begin{aligned}\bar{\mathfrak{g}}^T(x, z_x^T) &= \bar{G}_{i-1}^T(A_i^{T+1}) \\ &= TB_i^T + (Tb - (b+1)i)A_i^{T+1} \\ &= b(T-i)A_i^{T+1},\end{aligned}$$

where the final equality again follows from straightforward algebraic manipulations. This completes the proof of left-continuity at z_x^T . The proof of right-continuity in this case follows identically to our previous proof of right-continuity at A_j^{T+1} for all $j \in [0, T-1]$ s.t. $A_j^{T+1} > z_x^T$, and we omit the details. Combining the above completes the proof of the lemma. \square

6.7. Proof of Lemmas 9 and 10

Proof: [Proof of Lemma 9] The statements regarding continuity follow from Lemma 8. Noting that $d > \alpha_x^T$ implies $d > x$, concavity on (α_x^T, U) follows from (3.8) and Lemma 8. Let $i = \zeta_x^{T-1}$. Convexity on $(0, z_x^{T-1})$ follows from (3.8) and Lemma 8. Supposing $z_x^{T-1} \notin \{0, U\}$, it follows from definitions that $z_x^{T-1} \in (A_{i-1}^T, A_i^T]$, where $\mathcal{A}_x^T = A_i^T$. Combining with the convexity of $\bar{F}_j^{T-1}(x, d)$ and $\bar{G}_j^{T-1}(d)$ for all j , to prove the lemma, it suffices to demonstrate that: 1. $z_x^{T-1} \notin \bigcup_{j=-1}^{T-1} \{A_j^T\}$ implies that

$$\lim_{d \uparrow z_x^{T-1}} \partial_d^+ \mathfrak{f}^T(x, d) \leq \partial_d^+ \mathfrak{f}^T(x, z_x^{T-1}); \quad (6.57)$$

and 2. $x \notin \bigcup_{j=-1}^{T-1} \{A_j^T\}$ implies

$$\lim_{d \uparrow x} \partial_d^+ \mathfrak{f}^T(x, d) \leq \partial_d^+ \mathfrak{f}^T(x, x). \quad (6.58)$$

We treat several cases, and assume throughout that $z_x^{T-1}, x \notin \bigcup_{j=-1}^{T-1} \{A_j^T\}$. First, suppose $z_x^{T-1} = x$. In this case, it follows from a straightforward contradiction argument that $\Gamma_x^T = i = 0$, $x \in (0, A_0^T)$, and

$$\begin{aligned}\lim_{d \uparrow z_x^{T-1}} \partial_d^+ \mathfrak{f}^T(x, d) &= -1 + \lim_{d \uparrow z_x^{T-1}} \partial_d^+ \bar{F}_0^{T-1}(x, d) \\ &= -1 + (b-1)(T-1) - \frac{b+T-1}{A_0^{T-1}}x + 2\frac{b+T-1}{A_0^{T-1}}z_x^{T-1} \\ &= -1 + (b-1)(T-1) + \frac{b+T-1}{A_0^{T-1}}x \\ &\leq -1 + (b-1)(T-1) + \frac{b+T-1}{A_0^{T-1}}A_0^T = bT - (b+1).\end{aligned}$$

Alternatively,

$$\begin{aligned}\partial_d^+ \mathfrak{f}^T(x, z_x^{T-1}) &= b + \partial_d^+ \bar{G}_{-1}^{T-1}(x) \\ &= b + (T-1)b = bT.\end{aligned}$$

Combining the above completes the proof of both (6.57) and (6.58) in this case.

Next, suppose $z_x^{T-1} < x$. We again consider two cases, depending on how $\eta_d \triangleq x - d$ comes to go from lying above β_d^{T-1} to lying below β_d^{T-1} . In the case that the two actually intersect, it may be easily verified that $z_x^{T-1} = x - B_i^{T-1}$, and

$$\begin{aligned}\lim_{d \uparrow z_x^{T-1}} \partial_d^+ \mathfrak{f}^T(x, d) &= -1 + \lim_{d \uparrow z_x^{T-1}} \partial_d^+ \bar{F}_i^{T-1}(x, d) \\ &= -1 + (b-1)(T-1) - bi - \frac{b+T-1}{A_i^{T-1}}(z_x^{T-1} + B_i^{T-1}) + 2\frac{b+T-1}{A_i^{T-1}}z_x^{T-1} \\ &= -1 + (b-1)(T-1) - (b+1)i + \frac{b+T-1}{A_i^{T-1}}z_x^{T-1}.\end{aligned}$$

Alternatively,

$$\begin{aligned}\partial_d^+ \mathfrak{f}^T(x, z_x^{T-1}) &= -1 + \partial_d^+ \overline{G}_{i-1}^{T-1}(z_x^{T-1}) \\ &= -1 + (T-1)b - (b+1)i.\end{aligned}$$

It follows that

$$\partial_d^+ \mathfrak{f}^T(x, z_x^{T-1}) - \lim_{d \uparrow z_x^{T-1}} \partial_d^+ \mathfrak{f}^T(x, d) = T-1 - \frac{b+T-1}{A_i^{T-1}} z_x^{T-1},$$

which will be the same sign as

$$A_i^T - z_x^{T-1} \geq A_i^T - A_i^T = 0.$$

Combining the above completes the proof of (6.57) in this case. Furthermore, the proof of (6.58) follows from the fact that $x \notin \bigcup_{j=-1}^{T-1} \{A_j^T\}$, from which it follows that

$$\partial_d^+ \mathfrak{f}^T(x, x) - \lim_{d \uparrow x} \partial_d^+ \mathfrak{f}^T(x, d) = b+1.$$

Finally, suppose that $z_x^{T-1} < x$ and $z_x^{T-1} \notin \bigcup_{j=1}^{T-2} \{x - B_j^{T-1}\}$. As $x \notin \bigcup_{j=-1}^{T-1} \{A_j^T\}$, this final case follows nearly identically to the proof of (6.58) in the previous case, and we omit the details. Combining all of the above cases completes the proof of the lemma. \square

Proof: [Proof of Lemma 10] Note that for any measure $Q \in \mathfrak{M}(\mu)$, it is true that

$$\mathbb{E}_Q[f(D)] \leq \mathbb{E}_Q[\eta(D)] = \frac{f(R) - f(L)}{R - L} \mu + \frac{Rf(L) - Lf(R)}{R - L}.$$

But since q has support only on points $d \in [0, U]$ s.t. $f(d) = \eta(d)$, it follows that

$$\mathbb{E}_q[f(D)] = \mathbb{E}_q[\eta(D)] = \frac{f(R) - f(L)}{R - L} \mu + \frac{Rf(L) - Lf(R)}{R - L}.$$

Combining the above completes the proof. \square

6.8. Proof of Lemma 11

Proof: [Proof of Lemma 11] We first prove that $\mathcal{K}^T(x, d) \geq \mathfrak{f}^T(x, d)$ for all $d \in [0, U]$. Let $i = \zeta_x^{T-1}$, $j = \Gamma_x^{T-1}$, and $k = \Upsilon_x^{T-1}$. That $\mathcal{K}^T(x, 0) \geq \mathfrak{f}^T(x, 0)$ follows from definitions. We now prove that $\mathcal{K}^T(x, A_l^T) \geq \mathfrak{f}^T(x, A_l^T)$ for all $l \in [i, j-1]$. First, it will be useful to rewrite \mathcal{K} in a more convenient form. Noting that $\mathfrak{N}_x^T = A_k^T$,

$$\begin{aligned}\mathfrak{f}^T(x, A_k^T) &= b(A_k^T - x) + \overline{G}_{k-1}^{T-1}(A_k^T) \\ &= b(A_k^T - x) + (T-1)B_k^{T-1} + ((T-1)b - (b+1)k)A_k^T,\end{aligned}$$

and

$$\begin{aligned}\mathcal{K}^T(x, d) &= \frac{\mathfrak{f}^T(x, A_k^T) - Tx}{A_k^T} d + Tx \\ &= (Tb - (b+1)k + \frac{(T-1)B_k^{T-1} - (b+T)x}{A_k^T}) d + Tx \\ &= (b(T-k) - \frac{(b+T)x}{A_k^T}) d + Tx.\end{aligned}\tag{6.59}$$

Combining with the fact that for all $l \in [i, j-1]$ one has $A_l^T \in [z_x^{T-1}, x]$, proving the desired statement is equivalent to proving that

$$(b(T-k) - \frac{(b+T)x}{A_k^T})A_l^T + Tx \geq x - A_l^T + (T-1)B_l^{T-1} + ((T-1)b - (b+1)l)A_l^T,$$

which is itself equivalent to demonstrating that

$$(b(l+1-k) + 1 - \frac{(b+T)x}{A_k^T})A_l^T + (T-1)x \geq 0. \quad (6.60)$$

First, it will be useful to prove that the left-hand side of (6.60),

$$\eta(l) \triangleq (b(l+1-k) + 1 - \frac{(b+T)x}{A_k^T})A_l^T + (T-1)x,$$

is decreasing in l , for $l \in [0, j-1]$. Indeed, after simplifying, we find that

$$\eta(l+1) - \eta(l) = \left(b + \frac{b}{b+l+1} (b(l+1-k) + 1 - \frac{(b+T)x}{A_k^T}) \right) A_{l+1}^T,$$

which will be the same sign as

$$\begin{aligned} & (b(l+2-k) + l+2)A_k^T - (b+T)x \\ & \leq (b(l+2-k) + l+2)A_k^T - (b+T)B_k^T \\ & = (b+1)(l+2-k)A_k^T. \end{aligned}$$

Noting that $l+1 \leq j-1$ and (3.10) implies $j \leq k$ completes the proof of monotonicity, which we now use to complete the proof of (6.60). In particular, the above monotonicity implies that to prove (6.60), it suffices to prove that $\eta(k-1) \geq 0$. Note that

$$\eta(k-1) = A_{k-1}^T + \frac{b(T-1-k) - k}{b+k}x.$$

If $b(T-1-k) - k \geq 0$, then trivially $\eta(k-1) \geq 0$. Thus suppose $b(T-1-k) - k < 0$. In this case, $\eta(k-1) \geq 0$ iff

$$x \leq \frac{(b+k)A_{k-1}^T}{k - b(T-1-k)}.$$

As $x \leq B_{k+1}^T$, it thus suffices to prove that

$$B_{k+1}^T \leq \frac{b+k}{k - b(T-1-k)}A_{k-1}^T,$$

which, dividing both sides by $A_{k-1}^T(b+k)$ and simplifying, is itself equivalent to proving that

$$(b+T)k \geq (k - b(T-1-k))(b+k+1).$$

Noting that $k \leq T-1$, and thus $k - b(T-1-k) \leq k$, thus completes the desired proof that $\mathcal{K}^T(x, A_l^T) \geq \mathfrak{f}^T(x, A_l^T)$ for all $l \in [i, j-1]$.

We now prove that $\mathcal{K}^T(x, A_l^T) \geq \mathfrak{f}^T(x, A_l^T)$ for all $l \in [j, k]$. By construction,

$$\mathcal{K}^T(x, A_k^T) = \mathfrak{f}^T(x, A_k^T),$$

and thus it suffices to prove the desired claim for $l \in [j, k-1]$. Note that the degenerate case for which $x = A_k^T$ can be ignored, as in that case $j = k$. Thus suppose $x < A_k^T$. In this case, Lemma 9

implies that $\mathfrak{f}^T(x, d)$ is a continuous, concave, piecewise linear function of d on $[A_j^T, A_k^T]$. As $\mathcal{K}^T(x, d)$ is a linear function of d , it follows from the basic properties of concave functions that to prove the desired claim, it suffices to demonstrate that

$$\partial_d^+ \mathcal{K}(x, A_k^T) \leq \lim_{d \uparrow A_k^T} \partial_d^+ \mathfrak{f}^T(x, d), \quad (6.61)$$

which is equivalent to proving that

$$b(T-k) - \frac{(b+T)x}{A_k^T} \geq b + \lim_{d \uparrow A_k^T} \partial_d^+ \overline{G}_{k-1}^{T-1}(d). \quad (6.62)$$

It follows from definitions that

$$A_k^T = \frac{b+T}{k} B_k^T \leq \frac{b+T}{k} x.$$

Combining with (6.62), we find that to prove the desired claim, it suffices to demonstrate that

$$b(T-k) - k \geq b + (T-1)b - (b+1)k. \quad (6.63)$$

Noting that both sides of (6.63) are equivalent completes the proof.

Finally, let us prove that $\mathcal{K}^T(x, A_l^T) \geq \mathfrak{f}^T(x, A_l^T)$ for all $l \in [k+1, T-1]$, which will complete the proof. It again follows from Lemma 9 and the basic properties of concave functions that in this case it suffices to demonstrate that $\partial_d^+ \mathcal{K}^T(x, A_k^T) \geq \partial_d^+ \mathfrak{f}^T(x, A_k^T)$, which is equivalent to proving that

$$b(T-k) - \frac{(b+T)x}{A_k^T} \geq b + \partial_d^+ \overline{G}_k^{T-1}(A_k^T). \quad (6.64)$$

It follows from definitions that

$$\begin{aligned} A_k^T &= \frac{k+1}{b+k+1} A_{k+1}^T \\ &= \frac{b+T}{b+k+1} B_{k+1}^T \geq \frac{b+T}{b+k+1} x. \end{aligned}$$

Combining with (6.64), we find that in this case it suffices to demonstrate that

$$b(T-k) - (b+k+1) \geq b + (T-1)b - (b+1)(k+1). \quad (6.65)$$

Noting that both sides of (6.65) are equivalent completes the proof. Combining all of the above with the piece-wise convexity guaranteed by Lemma 9 completes the proof that $\mathcal{K}^T(x, d) \geq \mathfrak{f}^T(x, d)$ for all $d \in [0, U]$.

We now prove that for all $l \in [k, T-2]$, $\mathcal{L}_l^T(x, d) \geq \mathfrak{f}^T(x, d)$ for all $d \in [0, U]$. Note that $\mathfrak{f}^T(x, d)$ is a concave function on $[A_k^T, U]$, and by construction $\mathcal{L}_l^T(x, d)$ is a line tangent to $\mathfrak{f}^T(x, d)$ at A_l^T . It follows from the basic properties of concave functions that

$$\mathcal{L}_l^T(x, d) \geq \mathfrak{f}^T(x, d) \quad \text{for all } d \in [A_k^T, U].$$

Combining those same properties with (6.61) and the basic properties of linear functions, it follows that $\mathcal{L}_l^T(x, d) \geq \mathcal{K}^T(x, d)$ for all $d \in [0, A_k^T]$. Combining all of the above completes the proof. \square

6.9. Proof of Theorems 5 and 6, and Observation 3

Proof: [Proof of Theorem 5] Let π_1 denote the base-stock policy which always orders up to 0, and π_2 the base-stock policy which always orders up to U . Note that if $x_0 = 0$, then for any $Q \in \mathcal{M}_{\text{GEN}}$, w.p.1 $C_t^{\pi_1} = bd_t$, and $C_t^{\pi_2} = U - d_t$, $t \in [1, T]$, which implies that $\text{Opt}_{\text{GEN}}^T(\mu, U, b) \leq \min\{Tb\mu, T(U - \mu)\} = \text{Opt}_{\text{IND}}^T(\mu, U, b)$. \square

Proof: [Proof of Theorem 6] From Theorems 1 and 3,

$$\frac{\text{Opt}_{\text{MAR}}^T(\mu, U, b)}{\text{Opt}_{\text{IND}}^T(\mu, U, b)} = \frac{G_{\Gamma_\mu^T}(\beta_\mu^T, \mu)}{\min\{Tb\mu, T(U - \mu)\}} = \frac{\Gamma_\mu^T A_{\Gamma_\mu^T}^{T+1} + (Tb - (b+1)\Gamma_\mu^T)\mu}{\min\{Tb\mu, T(U - \mu)\}}.$$

Combining the above with Lemma 12 completes the proof. \square

Proof: [Proof of Observation 3] The results of [115] (i.e. Theorem 1) imply that if $x_0 = U$ and $\frac{\mu}{U} \leq \frac{1}{b+1}$, then the dynamics of Problem 2.1 at optimality are as follows. The initial inventory level equals U . In each period, with probability $\frac{\mu}{U}$, the demand equals U , and with probability $1 - \frac{\mu}{U}$ equals 0. Up until the first time that the demand equals U , the inventory level will equal U , and at the end of each period a holding cost equal to U is incurred. The first time the demand equals U , the inventory level drops to 0, and in that period no cost is incurred. In all later periods, the inventory level is raised to 0, and if in that period the demand equals U , a cost of $b\mu$ is incurred - otherwise no cost is incurred. It follows that the value of Problem 2.1 equals

$$\sum_{k=1}^T \left(1 - \frac{\mu}{U}\right)^{k-1} \frac{\mu}{U} ((k-1)U + (T-k)b\mu) + \left(1 - \frac{\mu}{U}\right)^T UT,$$

which is itself equal to

$$b\mu T + \frac{U^2}{\mu} - (b+1)U + \left(1 - \frac{\mu}{U}\right)^T \left(1 + b - \frac{U}{\mu}\right)U.$$

Alternatively, it follows from Theorems 2 and 3, combined with the fact that $\mathbf{q}_{U,\mu}^T(U) = \frac{\mu}{U}$, $\mathbf{q}_{U,\mu}^T(0) = 1 - \frac{\mu}{U}$, and the martingale property, that the dynamics of Problem 2.4 at optimality are as follows. The initial inventory level equals U . With probability $\frac{\mu}{U}$, the demand in every period is U , and the inventory level is raised to U at the start of each period, so no cost is incurred over the entire time horizon. With probability $1 - \frac{\mu}{U}$, the demand in every period is 0, and the inventory level in every period is U , implying that a cost of TU is incurred over the time horizon. It follows that the value of Problem 2.4 equals $(1 - \frac{\mu}{U})TU = (U - \mu)T$. Combining the above with a straightforward limiting argument, and the fact that $\frac{\mu}{U} < \frac{1}{b+1}$ implies $\frac{U-\mu}{b\mu} > 1$, completes the proof. \square