Risk-Constrained Nonconvex Dynamic Resource Allocation has Zero Duality Gap

Dionysis Kalogerias
Department of EE – Yale

June 2022

Abstract

We show that risk-constrained dynamic resource allocation problems with general integrable nonconvex instantaneous service functions exhibit zero duality gap. We consider risk constraints which involve convex and positively homogeneous risk measures admitting dual representations with bounded risk envelopes, and are strictly more general than expectations. Beyond expectations, particular risk measures supported within our setting include the conditional value-at-risk, the mean-absolute deviation (including the non-monotone case), certain distributionally robust representations and more generally all real-valued coherent risk measures on the space $\mathcal{L}_1$. Our proof technique relies on risk duality in tandem with Uhl's weak extension of Lyapunov's convexity theorem for vector measures taking values in general Banach spaces.

Notation: Bold capital letters (such as $\mathbf{A}$), or calligraphic letters (such as $\mathcal{A}$), or sometimes plain capital letters (such as A) will denote finite-dimensional sets. Double stroke letters (such as $\mathbb{A}$) will denote infinite-dimensional sets, such as Banach spaces. Math script letters (such as $\mathscr{A}$) will denote σ-algebras. Boldsymbol letters (such as $\boldsymbol{A}$ or $\boldsymbol{a}$) will denote (random) vectors. The space of integrable functions from a measurable space $(\Omega, \mathscr{F})$ equipped with a finite measure $\mu : \mathscr{F} \to \mathbb{R}_+$ to a Banach space $\mathbb{A}$, with standard notation $\mathcal{L}_1(\Omega, \mathscr{F}, \mu; \mathbb{A})$ [Diestel and Uhl, 1977, Shapiro et al., 2014], is abbreviated as $\mathcal{L}_1(\mu, \mathbb{A})$, as it is also common practice [Diestel and Uhl, 1977]. The rest of the notation is standard.

1 Introduction and Problem Setting

On some arbitrary base probability space $(\Omega, \mathscr{F}, \mu)$, consider a random element $\boldsymbol{H} : \Omega \to \mathcal{H} \triangleq \mathbb{R}^{N_{\boldsymbol{H}}}$ with induced Borel measure $\mathrm{P} : \mathscr{B}(\mathcal{H}) \to [0, 1]$, modeling some random phenomenon, which we would like to optimally manage by allocating resources. In particular, we are interested in the class of nonconvex functional resource allocation problems formulated as

$$\boxed{\begin{array}{rl} \underset{\boldsymbol{x},\boldsymbol{p}(\cdot)}{\text{maximize}} & g^o(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{x} \le -\boldsymbol{\rho}(-\boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})) \\ & \boldsymbol{g}(\boldsymbol{x}) \ge \mathbf{0} \\ & (\boldsymbol{x}, \boldsymbol{p}) \in \mathcal{X} \times \Pi \end{array}}\,, \tag{RCP}$$

where, with $\mathbf{C} \triangleq \mathbb{R}^N$, $g^o : \mathbf{C} \to \mathbb{R}$ and $\boldsymbol{g} : \mathbf{C} \to \mathbb{R}^{N_{\boldsymbol{g}}}$ are concave fixed utility functions, $\boldsymbol{p} : \mathcal{H} \to \mathbf{R} \triangleq \mathbb{R}^{N_{\boldsymbol{p}}}$ is the resource allocation policy, $\boldsymbol{f} : \mathbf{R} \times \mathcal{H} \to \mathbf{C}$ is a generally nonconvex *instantaneous* performance level score, measuring the quality of a policy $\boldsymbol{p}$ at each realization $\boldsymbol{H}$ in $\mathcal{H}$ and such that

$\boldsymbol{f}(\cdot, \boldsymbol{H}) \in \mathcal{L}_1(\mathrm{P}, \mathbf{C})$ on Π, and where $\boldsymbol{\rho} : \mathcal{L}_1(\mathrm{P}, \mathbf{C}) \to \mathbf{C}$ is a finite-valued vector risk measure, which we assume that is *convex, lower semicontinuous and positively homogeneous* in every dimension (i.e., component-wise), with the standardized convention that, for each $i \in \mathbb{N}_N^+$,

$$\rho_i(\boldsymbol{Z}) = \rho_i(Z_i), \quad \text{for all } \boldsymbol{Z} \in \mathcal{L}_1(\mathrm{P}, \mathbf{C}).$$

As indicated in (RCP), performance risks are further restricted to the finite-dimensional set $\mathcal{X} \subseteq \mathbf{C}$ and resource allocations are further restricted to the infinite-dimensional set Π. More specialized (yet general enough) assumptions on the structure of (RCP) which enable the development of the results advocated in this paper will be discussed in due course.

As an example, if $i \in \mathbb{N}_N^+$ refers to a user of a wireless system, then the i-th entry of $\boldsymbol{\rho}$ may evaluate the risk associated with the service experienced by the i-th user; hereafter, $\boldsymbol{f}$ will also be called the *service function*. However, with the exception of very special scenarios, the service experienced by the i-th user *can* (and, in general, will) be dependent on the services experienced by the rest of the users in the system. This is due to each service f_i being a function of the coupling variables $\boldsymbol{p}(\boldsymbol{H})$ and $\boldsymbol{H}$, which are common to all users in the system. This general problem structure adheres to a large variety of resource allocation problems encountered in practice.

Problem (RCP) admits an equivalent and rather useful representation. By the duality theorem for coherent risk measures [Shapiro et al., 2014, Theorem 6.5], we have that for every risk measure ρ which is convex, lower semicontinuous and positively homogeneous (i.e., in the class of those appearing in problem (RCP)), it holds that

$$\rho(Z) = \sup_{\zeta \in \mathbb{A}} \langle \zeta, Z \rangle \triangleq \sup_{\zeta \in \mathbb{A}} \int \zeta(\boldsymbol{h}) Z(\boldsymbol{h}) \mathrm{dP}(\boldsymbol{h}), \quad \text{for all } Z \in \mathcal{L}_1(\mathrm{P}, \mathbb{R}),$$

where the *uncertainty set* $\mathbb{A}$ is the domain of the Legendre-Fenchel conjugate of ρ, and it holds that $\mathbb{A} \subseteq \mathcal{L}_\infty(\mathrm{P}, \mathbb{R})$ by functional duality. The set $\mathbb{A}$ is also sometimes called the *risk envelope* of ρ. In this paper, we will be focusing on risk envelopes which are subsets of $\{\zeta \in \mathcal{L}_\infty(\mathrm{P}, \mathbb{R}) | |\zeta(\cdot)| \leq \gamma, \mathrm{P}\text{-a.e.}\}$, where $\gamma > 0$ is some arbitrarily large but finite constant, but otherwise we make no further explicit assumptions. Of course, it follows that, for every $Z \in \mathcal{L}_1(\mathrm{P}, \mathbb{R})$,

$$-\rho(-Z) = \inf_{\zeta \in \mathbb{A}} \langle \zeta, Z \rangle.$$

Under this provisioning, our initial functional program may be equivalently expressed as

$$\boxed{\begin{array}{ll} \underset{\boldsymbol{x}, \boldsymbol{p}(\cdot)}{\text{maximize}} & g^o(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{x} \leq \inf_{\boldsymbol{\zeta} \in \mathbb{A}_\gamma^S} \mathbb{E}\{\boldsymbol{\zeta}(\boldsymbol{H}) \circ \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})\} \\ & \boldsymbol{g}(\boldsymbol{x}) \geq \boldsymbol{0} \\ & (\boldsymbol{x}, \boldsymbol{p}) \in \mathcal{X} \times \Pi \end{array}}, \tag{RCP-E}$$

where "$\circ$" denotes the Hadamard product, the infimum is understood in a component-wise manner, and where the service uncertainty set (i.e., risk envelope) $\mathbb{A}_\gamma^S$ is defined as the Cartesian product $\mathbb{A}_\gamma^S \triangleq \mathbb{A}_\gamma^1 \times \mathbb{A}_\gamma^2 \times \ldots \times \mathbb{A}_\gamma^N$, with each $\mathbb{A}_\gamma^i$ satisfying the inclusion

$$\mathbb{A}_\gamma^i \subseteq \{\zeta \in \mathcal{L}_\infty(\mathrm{P}, \mathbb{R}) | |\zeta(\cdot)| \leq \gamma, \mathrm{P}\text{-a.e.}\}, \quad \forall i \in \mathbb{N}_N^+.$$

In passing, it may be worth noting that (RCP-E) is equivalent to the semi-infinite functional program

$$\begin{array}{ll} \underset{\boldsymbol{x}, \boldsymbol{p}(\cdot)}{\text{maximize}} & g^o(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{x} \leq \mathbb{E}\{\boldsymbol{\zeta}(\boldsymbol{H}) \circ \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})\}, \quad \forall \boldsymbol{\zeta} \in \mathbb{A}_\gamma^S \\ & \boldsymbol{g}(\boldsymbol{x}) \geq \boldsymbol{0} \\ & (\boldsymbol{x}, \boldsymbol{p}) \in \mathcal{X} \times \Pi \end{array}, \tag{RCP-I}$$

illustrating the generality of the risk-constrained problem (RCP), as compared with its risk-neutral counterpart obtained by choosing $\mathbb{A}_\gamma^i = \{\zeta(\cdot) = 1, \text{P-a.e.}\}$ for $i \in \mathbb{N}_N^+$, which is of course equivalent with replacing the vector risk measure $\boldsymbol{\rho}$ by a vector of expectations.

Problem (RCP-E) (which is equivalent to (RCP) under the conditions outlined above) is very general and several cases of special interest can be formulated as particular instances. As discussed in the previous paragraph, those instances include the risk-neutral version of (RCP), where all risk measures are trivially replaced by expectations. A nontrivial case of particular value is that of the Conditional Value-at-Risk at level $\beta \in (0,1]$ (CVaR^β) defined as [Rockafellar and Uryasev, 2000, Shapiro et al., 2014]

$$\text{CVaR}^\beta(Z) = \inf_{t\in\mathbb{R}} t + \frac{1}{\beta}\mathbb{E}\{(Z-t)_+\}, \quad Z \in \mathcal{L}_1(\text{P}, \mathbb{R}),$$

for which (RCP) reduces to

$$\boxed{\begin{array}{ll} \underset{\boldsymbol{x},\boldsymbol{p}(\cdot)}{\text{maximize}} & g^o(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{x} \leq -\text{CVaR}^{\boldsymbol{\beta}}(-\boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})) \\ & \boldsymbol{g}(\boldsymbol{x}) \geq \boldsymbol{0} \\ & (\boldsymbol{x}, \boldsymbol{p}) \in \mathcal{X} \times \Pi \end{array}}, \tag{CVaR}$$

where $\boldsymbol{\beta} \triangleq [\beta_1\, \beta_2\, \dots\, \beta_N] \in (0,1]^N$ is a vector containing the CVaR levels associated with each entry of the service function $\boldsymbol{f}$. In this case, the equivalent formulation of (CVaR) in the form of (RCP-E) is valid by choosing $\gamma = \max_{i\in\mathbb{N}_N^+} 1/\beta_i$, and with corresponding risk envelopes given by

$$\mathbb{A}_\gamma^i = \{\zeta \in \mathcal{L}_\infty(\text{P}, \mathbb{R}) | \zeta(\cdot) \in [0, 1/\beta_i], \text{P-a.e., and } \mathbb{E}\{\zeta\} = 1\}, \quad \forall i \in \mathbb{N}_N^+.$$

Note that in the case of CVaR, the resulting functional program may also be stated as

$$\begin{array}{ll} \underset{\boldsymbol{x},\boldsymbol{p}(\cdot)}{\text{maximize}} & g^o(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{x} \leq \sup\limits_{\boldsymbol{t}\in\mathbf{C}} \boldsymbol{t} + \dfrac{1}{\boldsymbol{\beta}} \circ \mathbb{E}\{-(\boldsymbol{t} - \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H}))_+\} \\ & \boldsymbol{g}(\boldsymbol{x}) \geq \boldsymbol{0} \\ & (\boldsymbol{x}, \boldsymbol{p}) \in \mathcal{X} \times \Pi \end{array},$$

which is the same as

$$\begin{array}{ll} \underset{\boldsymbol{x},\boldsymbol{p}(\cdot),\boldsymbol{t}\in\mathbf{C}}{\text{maximize}} & g^o(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{x} \leq \boldsymbol{t} + \dfrac{1}{\boldsymbol{\beta}} \circ \mathbb{E}\{-(\boldsymbol{t} - \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H}))_+\} \\ & \boldsymbol{g}(\boldsymbol{x}) \geq \boldsymbol{0} \\ & (\boldsymbol{x}, \boldsymbol{p}) \in \mathcal{X} \times \Pi \end{array}.$$

Another popular special case is that of the mean absolute deviation with trade-off parameter $\lambda \geq 0$ (MAD^λ) defined as [Ogryczak and Ruszczyński, 1999, Ogryczak and Ruszczyński, 2002]

$$\text{MAD}^\lambda(Z) = \mathbb{E}\{Z\} + \lambda\mathbb{E}\{|Z - \mathbb{E}\{Z\}|\}, \quad Z \in \mathcal{L}_1(\text{P}, \mathbb{R}),$$

for which problem (RCP) reduces to

$$\boxed{\begin{array}{ll} \underset{\boldsymbol{x},\boldsymbol{p}(\cdot)}{\text{maximize}} & g^o(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{x} \leq -\text{MAD}^{\boldsymbol{\lambda}}(-\boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})) \\ & \boldsymbol{g}(\boldsymbol{x}) \geq \boldsymbol{0} \\ & (\boldsymbol{x}, \boldsymbol{p}) \in \mathcal{X} \times \Pi \end{array}}, \tag{MAD}$$

where $\boldsymbol{\lambda} \triangleq [\lambda_1\, \lambda_2\, \ldots\, \lambda_N] \in \mathbb{R}_+^N$ is another vector containing the MAD trade-offs associated with each entry of the service function $\boldsymbol{f}$. In this case, (MAD) can be equivalently written in the form of (RCP-E) by choosing $\gamma = \max_{i \in \mathbb{N}_N^+} 1 + 2\lambda_i$, and risk envelopes

$$\mathbb{A}_\gamma^i = \{\zeta \in \mathcal{L}_\infty(\mathrm{P}, \mathbb{R}) | \zeta = 1 + \zeta' - \mathbb{E}\{\zeta'\},\ \text{and}\ \|\zeta'\|_{\mathcal{L}_\infty} \leq \lambda_i\}, \quad \forall i \in \mathbb{N}_N^+.$$

Similar to the case of CVaR, it is easy to see that the MAD program takes the form

$$\begin{array}{ll} \underset{\boldsymbol{x}, \boldsymbol{p}(\cdot)}{\text{maximize}} & g^o(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{x} \leq \mathbb{E}\{\boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})\} - \boldsymbol{\lambda} \circ \mathbb{E}\{|\mathbb{E}\{\boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})\} - \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})|\} \\ & \boldsymbol{g}(\boldsymbol{x}) \geq \boldsymbol{0} \\ & (\boldsymbol{x}, \boldsymbol{p}) \in \mathcal{X} \times \Pi \end{array},$$

or, equivalently,

$$\begin{array}{ll} \underset{\boldsymbol{x}, \boldsymbol{p}(\cdot), \boldsymbol{t} \in \mathbf{C}}{\text{maximize}} & g^o(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{x} \leq \boldsymbol{t} - \boldsymbol{\lambda} \circ \mathbb{E}\{|\boldsymbol{t} - \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})|\} \\ & \boldsymbol{t} = \mathbb{E}\{\boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})\} \\ & \boldsymbol{g}(\boldsymbol{x}) \geq \boldsymbol{0} \\ & (\boldsymbol{x}, \boldsymbol{p}) \in \mathcal{X} \times \Pi \end{array}.$$

There are numerous other risk measures which are compatible with our assumptions on the vector risk measure $\boldsymbol{\rho}$. In fact, all coherent risk measures on the space $\mathcal{L}_1$ (relative to any choice of the probability measure P) are supported under the adopted framework. For any such risk measure ρ, it is well known that

$$\rho(Z) = \sup_{\zeta \in \mathbb{A}} \langle \zeta, Z \rangle = \sup_{\frac{\mathrm{dQ}}{\mathrm{dP}} \in \mathbb{A}} \mathbb{E}\left\{ Z(\boldsymbol{H}) \frac{\mathrm{dQ}}{\mathrm{dP}}(\boldsymbol{H}) \right\}, \quad \text{for all } Z \in \mathcal{L}_1(\mathrm{P}, \mathbb{R}),$$

where the risk envelope $\mathbb{A}$ takes the special form

$$\begin{aligned} \mathbb{A} &= \{\zeta \in \mathcal{L}_\infty(\mathrm{P}, \mathbb{R}) | \langle \zeta, Z \rangle \leq \rho(Z) \text{ for all } Z \in \mathcal{L}_1(\mathrm{P}, \mathbb{R}), \zeta \succeq 0, \mathbb{E}\{\zeta\} = 1\} \\ &= \left\{ \frac{\mathrm{dQ}}{\mathrm{dP}} \in \mathcal{L}_\infty(\mathrm{P}, \mathbb{R}) \middle| \left\langle \frac{\mathrm{dQ}}{\mathrm{dP}}, Z \right\rangle \leq \rho(Z),\ \text{for all } Z \in \mathcal{L}_1(\mathrm{P}, \mathbb{R}) \right\}. \end{aligned}$$

In the above, $\mathrm{dQ}/\mathrm{dP} = \zeta$ denotes the Radon-Nikodym derivative of a probability measure Q on $(\mathcal{H}, \mathscr{B}(\mathcal{H}))$ relative to P, with the former assumed absolutely continuous relative to the latter; we use the standard notation $\mathrm{Q} \ll \mathrm{P}$.

If ρ is also real-valued, it is known that it is continuous in the strong topology on $\mathcal{L}_1(\mathrm{P}, \mathbb{R})$, while its risk envelope $\mathbb{A}$ is a convex bounded and weakly*-closed subset of the set of densities in $\mathcal{L}_\infty(\mathrm{P}, \mathbb{R})$, also implying that the supremum is attained for some $\zeta_Z^\star \in \mathbb{A}$ [Shapiro et al., 2014, Proposition 6.6 and Theorem 6.7]. Since, in particular, $\mathbb{A}$ is bounded, there must exist a $\gamma > 0$ such that $|\zeta(\cdot)| \leq \gamma$, P-a.e. for every $\zeta \in \mathbb{A}$ (note that such a γ is necessarily independent of the choice of Z). Therefore, $\mathbb{A}$ is consistent with our setting. It follows that ρ admits a distributionally robust representation of the form

$$\rho(Z) = \sup_{\mathrm{Q} \in \mathfrak{M}} \mathbb{E}_{\mathrm{Q}}\{Z(\boldsymbol{H})\}, \quad \text{for all } Z \in \mathcal{L}_1(\mathrm{P}, \mathbb{R}),$$

where

$$\mathfrak{M} \triangleq \left\{ \mathrm{Q} \ll \mathrm{P} \middle| \left\langle \frac{\mathrm{dQ}}{\mathrm{dP}}, Z \right\rangle \leq \rho(Z),\ \text{for all } Z \in \mathcal{L}_1(\mathrm{P}, \mathbb{R}),\ \text{with } \frac{\mathrm{dQ}}{\mathrm{dP}}(\cdot) \leq \gamma, \mathrm{P}\text{-a.e.} \right\}$$

$$\subseteq \left\{ \mathrm{Q} \ll \mathrm{P} \middle| \mathrm{Q} \leq \gamma \mathrm{P}, \text{ on Borel sets} \right\}.$$

To further highlight the versatility of (RCP), would like to emphasize that one can freely "mix-and-match" any combination of risk measures present in the constraints (RCP), out of the large variety of those supported under our assumptions. Such combinations include different choices of risk measures across the components of the service vector $\boldsymbol{f}$, e.g., expectations for some components and CVaRs or MADs for others, as well as combinations of different risk measures for each of the components of $\boldsymbol{f}$; a classical example is mean-CVaR trade-offs.

2 Lagrangian Duality

A promising and celebrated approach for dealing with the explicit constraints of the risk-aware problem (RCP) is by exploiting *Lagrangian duality*, which has been proven essential in analyzing and efficiently solving constrained convex optimization problems; see, e.g., Boyd and Vandenberghe [2004], Ruszczyński [2006], Bertsekas [2009]. Note, however, that since in problem (RCP) the services in $\boldsymbol{f}(\cdot, \boldsymbol{H})$ are nonconcave in general (i.e., with respect to the first argument corresponding to the policy $\boldsymbol{p}$), most standard results in Lagrangian duality for convex optimization do not apply automatically. On top of that, one has to incorporate the structural complexity of the risk measure $\boldsymbol{\rho}$, which is in general a nonlinear functional of its argument. As a result, fundamental properties of expectation (being the most trivial risk measure) which enable an elegant and straightforward analysis, such as linearity, do not hold for $\boldsymbol{\rho}$.

The *Lagrangian function* $\mathsf{L} : \mathbf{C} \times \Pi \times \mathbb{R}^{N_{\boldsymbol{g}}} \times \mathbf{C} \to \mathbb{R}$ associated with the risk-constrained problem (RCP) is defined by scalarizing its constraints as

$$\mathsf{L}(\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\lambda}) \triangleq g^o(\boldsymbol{x}) + \langle \boldsymbol{\lambda}_{\boldsymbol{g}}, \boldsymbol{g}(\boldsymbol{x}) \rangle + \left\langle \boldsymbol{\lambda}_{\boldsymbol{\rho}}, -\boldsymbol{\rho}(-\boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})) - \boldsymbol{x} \right\rangle,$$

where $\boldsymbol{\lambda} \equiv (\boldsymbol{\lambda}_{\boldsymbol{g}}, \boldsymbol{\lambda}_{\boldsymbol{\rho}}) \in \mathbb{R}^{N_{\boldsymbol{g}}} \times \mathbf{C}$ are the multipliers associated with the constraints of (RCP). Then the *dual function* $\mathsf{D} : \mathbb{R}^{N_{\boldsymbol{g}}} \times \mathbf{C} \to (-\infty, \infty]$ is defined as

$$\mathsf{D}(\boldsymbol{\lambda}) \triangleq \sup_{(\boldsymbol{x}, \boldsymbol{p}) \in \mathcal{X} \times \Pi} \mathsf{L}(\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\lambda}).$$

If the optimal value of problem (RCP) is $\mathsf{P}^* \in \mathbb{R}$, it is easily understood that that $\mathsf{P}^* \leq \mathsf{D}$ on the positive orthant (i.e., for $\boldsymbol{\lambda} \geq \mathbf{0}$), and thus it is most reasonable to consider the *dual problem* $\inf_{\boldsymbol{\lambda} \geq \mathbf{0}} \mathsf{D}(\boldsymbol{\lambda})$, which is always convex and whose optimal value

$$\mathsf{D}^* \triangleq \inf_{\boldsymbol{\lambda} \geq \mathbf{0}} \mathsf{D}(\boldsymbol{\lambda}) \in (-\infty, \infty]$$

serves as the tightest over-estimate of the optimal value of problem (RCP), P^*, when knowing only D. Then, one of the basic questions in Lagrangian duality is whether we can essentially replace an original constrained problem with its dual, in the sense that $\mathsf{P}^* \equiv \mathsf{D}^*$; in such a case, we say that the problem has *zero duality gap*. Referring to problem (RCP), this would imply that it can be replaced by the *minimax problem*

$$\inf_{\boldsymbol{\lambda} \geq \mathbf{0}} \mathsf{D}(\boldsymbol{\lambda}) \equiv \inf_{\boldsymbol{\lambda} \geq \mathbf{0}} \sup_{(\boldsymbol{x}, \boldsymbol{p}) \in \mathcal{X} \times \Pi} \mathsf{L}(\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\lambda}), \tag{1}$$

whose optimal value is D^*. A zero duality gap also implies that problem (RCP) satisfies the *saddle point property* (whether a saddle point exists or not), which is expressed as

$$\mathsf{D}^* = \inf_{\boldsymbol{\lambda} \geq \mathbf{0}} \sup_{(\boldsymbol{x}, \boldsymbol{p}) \in \mathcal{X} \times \Pi} \mathsf{L}(\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\lambda}) = \sup_{(\boldsymbol{x}, \boldsymbol{p}) \in \mathcal{X} \times \Pi} \inf_{\boldsymbol{\lambda} \geq \mathbf{0}} \mathsf{L}(\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\lambda}) = \mathsf{P}^*.$$

Zero duality gaps are important: Because there is a finite number of constraints, the dual function is finite dimensional even though the original functional problem (RCP) is infinite dimensional. Additionally, for every choice of the dual variable $\boldsymbol{\lambda} \geq \mathbf{0}$ (and therefore any optimal multiplier vector as well), maximization of the Lagrangian $\mathsf{L}(\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\lambda})$ over the pair $(\boldsymbol{x}, \boldsymbol{p}) \in \mathcal{X} \times \Pi$ is separable. We thus see that duality transforms a constrained problem into an unconstrained problem in a principled and favorable way and, provided that the original constrained stochastic program (in our case (RCP)) exhibits zero duality gap, presents a general methodology to tackle it. While zero duality gaps are a common and fundamental characteristic of problems in convex (concave) constrained optimization (under some regularity conditions), proving zero duality gaps in nonconcave problems such as (RCP) is a much more delicate and challenging task.

3 Main Results

Our structural setting regarding problem (RCP) is mostly standard as compared with the relevant literature on risk-neutral stochastic resource allocation and related problems [Luo and Zhang, 2008, Ribeiro, 2012, Chamon et al., 2020, 2021], and is as follows.

Assumption 1 (Structural Setting). *The following conditions are in effect:*

1) *The utilities g^o and $\boldsymbol{g}$ are concave.*

2) *The service feasible set $\mathcal{X}$ is convex.*

3) *The policy feasible set Π is decomposable.*

4) *The Borel measure P is nonatomic*[a]*.*

5) *Problem (RCP) satisfies Slater's condition (i.e., it is strictly feasible).*

6) *The risk limit span $\$ \triangleq \left\{\boldsymbol{s} \in \mathbf{C} \middle| \boldsymbol{s} = -\boldsymbol{\rho}(-\boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})),\ \text{for some } \boldsymbol{p} \in \Pi\right\}$ is closed.*

[a] Recall that P is nonatomic if for any event E with $\mathrm{P}(E) > 0$, an event $E' \subseteq E$ exists such that $\mathrm{P}(E) > \mathrm{P}(E') > 0$.

Condition **(6)** is the only critical part of Assumption 1 that does not appear in the state of the art [Luo and Zhang, 2008, Ribeiro, 2012, Chamon et al., 2020, 2021], and essentially means that *every possible risk can be realized by a feasible policy.* Since condition **(6)** is new, it is crucial that we can demonstrate its compatibility with the rest of the conditions of Assumption 1. In this respect, we first have the following result.

Lemma 1 (Span Compactness). *Suppose that the set Π is connected, that the composite functional $\boldsymbol{\rho}(-\boldsymbol{f}(\cdot, \boldsymbol{H}))$ is continuous in the strong (norm) topology on Π, and that the component-wise extrema of $\boldsymbol{\rho}(-\boldsymbol{f}(\cdot, \boldsymbol{H}))$ are attained on Π. Then, the risk limit span $\$$ is compact in $\mathbf{C}$.*

Proof of Lemma 1. Fix $i \in \mathbb{N}_N^+$ and suppose that, for every $\boldsymbol{p} \in \Pi$,

$$-\rho_i(-f_i(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})) \leq \sup_{\boldsymbol{p} \in \Pi} -\rho_i(-f_i(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})) = -\rho_i(-f_i(\boldsymbol{p}_u^i(\boldsymbol{H}), \boldsymbol{H})) < \infty \quad \text{and}$$
$$-\rho_i(-f_i(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})) \geq \inf_{\boldsymbol{p} \in \Pi} -\rho_i(-f_i(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})) = -\rho_i(-f_i(\boldsymbol{p}_l^i(\boldsymbol{H}), \boldsymbol{H})) > -\infty.$$

Because $-\rho_i(-f_i(\cdot, \boldsymbol{H}))$ is continuous in the strong topology on Π, and since Π is connected, it follows that the range (i.e., image) of $-\rho_i(-f_i(\cdot, \boldsymbol{H}))$ on Π is the connected, closed and bounded, and therefore compact interval

$$[-\rho_i(-f_i(\boldsymbol{p}_l^i(\boldsymbol{H}),\boldsymbol{H})), -\rho_i(-f_i(\boldsymbol{p}_u^i(\boldsymbol{H}),\boldsymbol{H}))].$$

This holds for every $i \in \mathbb{N}_N^+$, implying that the risk limit span (i.e., the range/image)

$$\mathbb{S} = \times_{i\in\mathbb{N}_N^+}[-\rho_i(-f_i(\boldsymbol{p}_l^i(\boldsymbol{H}),\boldsymbol{H})), -\rho_i(-f_i(\boldsymbol{p}_u^i(\boldsymbol{H}),\boldsymbol{H}))]$$

is also compact in $\mathbb{R}^N = \mathbf{C}$. ■

The strongest assumption of Lemma 1 is that $\boldsymbol{\rho}(-\boldsymbol{f}(\cdot,\boldsymbol{H}))$ attains its extrema component-wise on the space of feasible policies Π. Nevertheless, this assumption is natural and can be verified on a case-by-case basis for a variety of interesting problems. More generally, if we assume that $\boldsymbol{\rho}$ is monotone in every dimension, and that for each $i \in \mathbb{N}_N^+$ there exist deterministic vectors $\boldsymbol{p}_m^i$, $\boldsymbol{p}_M^i$ in $\mathbf{R}$ *and* in Π such that

$$m_i(\boldsymbol{H}) \triangleq f_i(\boldsymbol{p}_m^i,\boldsymbol{H}) \le f_i(\boldsymbol{p}(\boldsymbol{H}),\boldsymbol{H}) \le f_i(\boldsymbol{p}_M^i,\boldsymbol{H}) \triangleq M_i(\boldsymbol{H}),$$

then it is also the case that

$$-\rho_i(-m_i(\boldsymbol{H})) \le -\rho_i(-f_i(\boldsymbol{p}(\boldsymbol{H}),\boldsymbol{H})) \le -\rho_i(-M_i(\boldsymbol{H})),$$

where the respective upper and lower bounds are attained by feasible policies.

For instance, for a plethora of resource allocation problems arising in wireless communications it is true that $\boldsymbol{\rho}(-\boldsymbol{f}(\mathbf{0},\boldsymbol{H})) \equiv \mathbf{0}$, which in other words and according to our relevant example in Section 1 intuitively means that *allocation of zero resources at each user implies zero service risk*. On the other hand, it is also often reasonable that *maximal service level for each user is achieved by (feasibly) allocating all available resources (e.g., power) to that user, and zero resources to every other user*. Then, for the i-th user we have

$$0 \le -\rho_i(-f_i(\boldsymbol{p}(\boldsymbol{H}),\boldsymbol{H})) \le -\rho_i(-f_i(p_{Max}^i\boldsymbol{e}^i,\boldsymbol{H})),$$

where $\boldsymbol{e}^i$ denotes the i-th member of the standard Euclidean basis on $\mathbf{R}$. It is evident that both lower and upper bounds are attained by feasible policies for all users in the wireless system, assuming of course that deterministic policies are feasible. A standard choice of a decomposable and connected feasible set Π which also includes deterministic policies is the uniform box

$$\Pi = \left\{\boldsymbol{p} : \mathcal{H} \to \mathbf{R} \middle| \operatorname{ess\,sup}_{\mathrm{P}} \|\boldsymbol{p}(\cdot)\|_\infty \le U\right\},$$

where $U > 0$ is an appropriate fixed number, or, the more refined rectangular box

$$\Pi = \left\{\boldsymbol{p} : \mathcal{H} \to \mathbf{R} \middle| \operatorname{ess\,sup}_{\mathrm{P}} p^i(\cdot) \le U^i, i \in \mathbb{N}_{N_{\boldsymbol{p}}}^+\right\},$$

where the U^i's are fixed. In a more delicate setting, another pretty standard choice for Π is (see, e.g., [Ribeiro, 2012] and references therein)

$$\Pi = \left\{\boldsymbol{p} : \mathcal{H} \to \mathbf{R} \middle| \boldsymbol{p}(\boldsymbol{h}) \in \mathcal{U}(\boldsymbol{h}), \text{P-a.e.}\right\},$$

where $\mathcal{U} : \mathcal{H} \rightrightarrows \mathbf{R}$ is a compact-valued (measurable) multifunction, which is *also* closed [Shapiro et al., 2014, Section 7.2.3]. In other words, every feasible $\boldsymbol{p} \in \Pi$ is constrained to be a well-behaved Borel-measurable selection of $\mathcal{U}$.

Lastly, norm continuity of the composition $\boldsymbol{\rho}(-\boldsymbol{f}(\cdot,\boldsymbol{H}))$, as required by Lemma 1, can be guaranteed "in stages", as the next result suggests.

Lemma 2 (**Continuity in Stages**). *Let $p, q \in [1,\infty]$ be such that $1/p + 1/q = 1$. If $\boldsymbol{\rho}$ is component-wise continuous in the strong topology on $\mathcal{L}_1(\mathrm{P},\mathbb{R})$, Π is a subset of $\mathcal{L}_p(\mathrm{P},\mathbf{R})$ and $\boldsymbol{f}(\cdot,\boldsymbol{h})$ is Lipschitz on the corresponding decision space for every $\boldsymbol{h} \in \mathcal{H}$ with constant $L(\boldsymbol{h}) \in \mathcal{L}_q(\mathrm{P},\mathbb{R})$, then $\boldsymbol{\rho}(-\boldsymbol{f}(\cdot,\boldsymbol{H}))$ is continuous in the strong topology on Π.*

Proof of Lemma 2. For every pair of policies $\boldsymbol{p} \in \Pi$ and $\boldsymbol{p}' \in \Pi$, it is true that (note that any combination of Euclidean ℓ_p-norms on $\mathbf{C}$ and $\mathbf{R}$ will work just fine)

$$\begin{aligned}\|\boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}),\boldsymbol{H}) - \boldsymbol{f}(\boldsymbol{p}'(\boldsymbol{H}),\boldsymbol{H})\|_{\mathcal{L}_1} &\le \big\| L(\boldsymbol{H}) \|\boldsymbol{p}(\boldsymbol{H}) - \boldsymbol{p}'(\boldsymbol{H})\|_\infty \big\|_{\mathcal{L}_1} \\ &\le \big\| L(\boldsymbol{H})\big\|_{\mathcal{L}_q} \big\| \|\boldsymbol{p}(\boldsymbol{H}) - \boldsymbol{p}'(\boldsymbol{H})\|_\infty \big\|_{\mathcal{L}_p} \\ &= \big\| L(\boldsymbol{H})\big\|_{\mathcal{L}_q} \big\|\boldsymbol{p}(\boldsymbol{H}) - \boldsymbol{p}'(\boldsymbol{H})\big\|_{\mathcal{L}_p}.\end{aligned}$$

This shows that the function $\boldsymbol{f}(\cdot,\boldsymbol{H})$ from Π to $\mathcal{L}_1(\mathrm{P},\mathbf{C})$ is (strongly) continuous. For each $i \in \mathbb{N}_N^+$, ρ_i is also (strongly) continuous on $\mathcal{L}_1(\mathrm{P},\mathbb{R})$. As a result, the functional $\rho_i(-f_i(\cdot,\boldsymbol{H}))$ is (strongly) continuous on Π, which implies that $\boldsymbol{\rho}(-\boldsymbol{f}(\cdot,\boldsymbol{H}))$ is also continuous on Π in the strong topology. ■

Note that most popular risk measures, such as the CVaR (for $\beta \in (0,1]$) and MAD (for $\lambda \ge 0$) are continuous, as Lemma 2 demands. More generally, every real-valued convex and monotone risk measure is continuous [Shapiro et al., 2014, Proposition 6.6], and hence so is every coherent risk measure. For those and several other cases of continuous risk measures, the reader is referred to [Shapiro et al., 2014]. Additionally, the assumption that the service vector $\boldsymbol{f}(\cdot,\boldsymbol{h})$ is Lipschitz for every $\boldsymbol{h} \in \mathcal{H}$ with an integrable constant is also standard and reasonable, and can be readily verified in many practical cases.

Regarding the rest of the conditions of Assumption 1, all are standard as discussed above. We would just like to point out that condition **(4)** holds naturally if the Borel measure P has a density with respect with the Lebesgue measure; this is a valid assumption in numerous practical settings.

We are now in position to state the main result of this work. The detailed proof is presented next in Section 4.

Theorem 1 (**Strong Duality of Risk-Constrained Resource Allocation**). *Let Assumption 1 be in effect. Then problem (RCP) has zero duality gap, i.e.,* $\mathsf{P}^* = \mathsf{D}^*$.

4 Proof of Theorem 1

Hereafter, we let Assumption 1 be in effect. We will work with the utility-constraint set associated with problem (RCP) (or equivalently with problem (RCP-E))

$$\mathcal{C} \triangleq \left\{ (\delta_o, \boldsymbol{\delta}_r, \boldsymbol{\delta}_d) \;\middle|\; \begin{array}{l} g^o(\boldsymbol{x}) \ge \delta_o \\ -\boldsymbol{\rho}(-\boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}),\boldsymbol{H})) - \boldsymbol{x} \ge \boldsymbol{\delta}_r\,,\ \text{for some } (\boldsymbol{x},\boldsymbol{p}) \in \mathcal{X} \times \Pi \\ \boldsymbol{g}(\boldsymbol{x}) \ge \boldsymbol{\delta}_d \end{array} \right\}.$$

Following [Chamon et al., 2020, 2021], showing that $\mathcal{C}$ is convex and the strict feasibility of problem (RCP) (from Assumption 1) suffice to ensure that (RCP) exhibits zero duality gap, as a relatively simple consequence of the supporting hyperplane theorem; see Section 4.3, or [Chamon et al., 2020, Theorem 1 and its proof], or [Chamon et al., 2021, Appendix A] for the details. Proving convexity of

$\mathcal{C}$ is nontrivial in the case of (RCP) though, and does not follow from the analyses presented in the aforementioned articles. This is due to the nonlinearity of the functionals present in the risk constraints of (RCP), in sharp contrast to standard problems considered in the literature (e.g., in [Luo and Zhang, 2008, Ribeiro, 2012, Chamon et al., 2020, 2021]), where the corresponding constraints evaluate the service vector $\boldsymbol{f}$ solely through linear functionals, or in other words, expectations.

We should therefore prove that $\mathcal{C}$ is convex. This means that if $(\delta_o, \boldsymbol{\delta}_r, \boldsymbol{\delta}_d) \in \mathcal{C}$ for $(\boldsymbol{x}, \boldsymbol{p}) \in \mathcal{X} \times \Pi$ and if $(\delta'_o, \boldsymbol{\delta}'_r, \boldsymbol{\delta}'_d) \in \mathcal{C}$ for $(\boldsymbol{x}', \boldsymbol{p}') \in \mathcal{X} \times \Pi$, then, for every $\alpha \in [0,1]$, it should be the case that

$$\alpha(\delta_o, \boldsymbol{\delta}_r, \boldsymbol{\delta}_d) + (1-\alpha)(\delta'_o, \boldsymbol{\delta}'_r, \boldsymbol{\delta}'_d) \in \mathcal{C}.$$

In other words, we have to show that there exists another feasible pair $(\boldsymbol{x}_\alpha, \boldsymbol{p}_\alpha) \in \mathcal{X} \times \Pi$, such that

$$g^o(\boldsymbol{x}_\alpha) \geq \alpha\delta_o + (1-\alpha)\delta'_o \tag{2}$$

$$-\boldsymbol{\rho}(-\boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})) - \boldsymbol{x}_\alpha \geq \alpha\boldsymbol{\delta}_r + (1-\alpha)\boldsymbol{\delta}'_r \quad \text{and} \tag{3}$$

$$\boldsymbol{g}(\boldsymbol{x}_\alpha) \geq \alpha\boldsymbol{\delta}_d + (1-\alpha)\boldsymbol{\delta}'_d \tag{4}$$

The proof will be complete if we show that, for every $\alpha \in [0,1]$, there is a feasible policy $\boldsymbol{p}_\alpha \in \Pi$, such that

$$\boxed{-\boldsymbol{\rho}(-\boldsymbol{f}(\boldsymbol{p}_\alpha(\boldsymbol{H}), \boldsymbol{H})) \geq -\alpha\boldsymbol{\rho}(-\boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})) - (1-\alpha)\boldsymbol{\rho}(-\boldsymbol{f}(\boldsymbol{p}'(\boldsymbol{H}), \boldsymbol{H})).} \tag{5}$$

To see this, suppose that (5) is true, and let $\boldsymbol{x}_\alpha = \alpha\boldsymbol{x} + (1-\alpha)\boldsymbol{x}' \in \mathcal{X}$ (by assumption, $\mathcal{X}$ is convex). Then, concavity of g^o readily implies that

$$\begin{aligned} g^o(\boldsymbol{x}_\alpha) &= g^o(\alpha\boldsymbol{x} + (1-\alpha)\boldsymbol{x}') \\ &\geq \alpha g^o(\boldsymbol{x}) + (1-\alpha)g^o(\boldsymbol{x}') \\ &\geq \alpha\delta_o + (1-\alpha)\delta'_o, \end{aligned}$$

which verifies (2). Similarly we obtain validity of (4), and (3) using (5). Therefore, in the rest of the proof, we will work to show that (5) is indeed true. In passing, note that (5) is trivially true in the special case where $\boldsymbol{f}(\cdot, \boldsymbol{h})$ is concave for every $\boldsymbol{h} \in \mathcal{H}$ and Π is also a convex set, simply by choosing $\boldsymbol{p}_\alpha$ as a convex combination of $\boldsymbol{p}$ and $\boldsymbol{p}'$, and exploiting convexity and positive homogeneity of the vector risk measure $\boldsymbol{\rho}$.

4.1 Preliminaries on Vector Measures in Banach Spaces

Let us first introduce some basic definitions and notation from the study of vector measures taking values in general, infinite dimensional Banach spaces. For a comprehensive treatment of the subject, the reader is referred to the classical monograph [Diestel and Uhl, 1977].

Suppose that $\mathbb{X}$ is a possibly infinite-dimensional Banach space. A *vector measure* on a measurable space $(\Omega, \mathscr{F})$ is a function $\boldsymbol{G} : \mathscr{F} \to \mathbb{X}$. A vector measure $\boldsymbol{G}$ is called countably additive in the same fashion as a regular real-valued measure. The *variation* of a vector measure $\boldsymbol{G}$ is another function on sets $|\boldsymbol{G}| : \mathscr{F} \to \mathbb{R}_+$ defined as

$$|\boldsymbol{G}|(E) \triangleq \sup_{\pi \text{ is a finite partition of } E} \sum_{A \in \pi} \|\boldsymbol{G}(A)\|_{\mathbb{X}}.$$

Then $\boldsymbol{G}$ is suggestively said to be of *bounded variation* whenever $|\boldsymbol{G}|(\Omega) < \infty$. A vector measure $\boldsymbol{G}$ is called *nonatomic* if every event $E \in \mathscr{F}$ such that $\boldsymbol{G}(E) \neq 0$ can be partitioned into events E'

and $E \backslash E'$ such that $\boldsymbol{G}(E') \neq 0$ and $\boldsymbol{G}(E \backslash E') \neq 0$; in other words, every event of non-zero measure can be partitioned into two events of non-zero measure.

Further, in the following we will be using the concept of a *Bochner integral*, which is a now standard extension of the Lebesgue integral for functions taking values in infinite-dimensional Banach spaces. While we do not provide a formal description here, the reader is referred to the excellent exposition in [Diestel and Uhl, 1977, Section II], which is the standard textbook on the subject.

Lastly, our analysis will be based on the following extension to the celebrated Convexity Theorem of Lyapunov, due to Uhl [Uhl, 1969, Theorem 1 and last paragraph before References section], see also [Diestel and Uhl, 1977, Theorem IX.1.10]. This also classical result conveniently generalizes the convexity theorem to infinite-dimensional Banach spaces, albeit with some nontrivial provisioning on the topological properties of the range of the involved vector measure.

Theorem 2 (**[Uhl, 1969] Weak Lyapunov Theorem for the Strong Topology**). *Let $(\Omega, \mathscr{F})$ be a measurable space, and let $\mathbb{X}$ be any Banach space. Let $\boldsymbol{G} : \mathscr{F} \to \mathbb{X}$ be a countably additive vector measure of bounded variation. If $\boldsymbol{G}$ is nonatomic and admits a Radon-Nikodym representation, i.e., there exist a finite measure $\mu : \mathscr{F} \to \mathbb{R}_+$ and a function $\boldsymbol{f} \in \mathcal{L}_1(\mu, \mathbb{X})$ such that*

$$\boldsymbol{G}(E) = \int_E \boldsymbol{f}(\omega) \mathrm{d}\mu(\omega), \quad E \in \mathscr{F},$$

then the norm-closure of the range $\boldsymbol{G}(\mathscr{F})$ is convex and norm-compact.

4.2 Core of the Proof: Convexity of $\mathcal{C}$

We identify with $\mathbb{X}$ the Banach space of all real-valued sequences bounded in the sup norm, i.e., $\mathbb{X} = \ell_\infty$. For a feasible policy $\boldsymbol{p} \in \Pi$, define the vector measure $\boldsymbol{G}_{\boldsymbol{p}} : \mathscr{B}(\mathcal{H}) \to \mathbb{X}$ as

$$\boldsymbol{G}_{\boldsymbol{p}}(E) \triangleq \begin{bmatrix} \int_E \lambda_0(\boldsymbol{h}) \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{h}), \boldsymbol{h}) \mathrm{dP}(\boldsymbol{h}) \\ \int_E \lambda_1(\boldsymbol{h}) \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{h}), \boldsymbol{h}) \mathrm{dP}(\boldsymbol{h}) \\ \vdots \\ \int_E \lambda_n(\boldsymbol{h}) \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{h}), \boldsymbol{h}) \mathrm{dP}(\boldsymbol{h}) \\ \vdots \end{bmatrix} = \begin{bmatrix} \mathbb{E}\{\mathbb{1}_E(\boldsymbol{H}) \lambda_0(\boldsymbol{H}) \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})\} \\ \mathbb{E}\{\mathbb{1}_E(\boldsymbol{H}) \lambda_1(\boldsymbol{H}) \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})\} \\ \vdots \\ \mathbb{E}\{\mathbb{1}_E(\boldsymbol{H}) \lambda_n(\boldsymbol{H}) \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})\} \\ \vdots \end{bmatrix}, \ E \in \mathscr{B}(\mathcal{H}),$$

where $\mathbb{B} \triangleq \{\lambda_n\}_{n \in \mathbb{N}}$ is a countable basis (i.e., dense) on $\mathcal{L}_1(\mathrm{P}, \mathbb{R})$. We will be later using the basis $\mathbb{B}$ in order to approximate elements in $\mathbb{A}_\gamma^i$, $i \in \mathbb{N}_N^+$, each of which is a bounded subset of $\mathcal{L}_\infty(\mathrm{P}, \mathbb{R})$, and thus of $\mathcal{L}_1(\mathrm{P}, \mathbb{R})$ (since P is finite); therefore, without loss of generality we may very well assume that $|\lambda_n| \leq \gamma$ everywhere on $\mathcal{H}$, and we do so hereafter (otherwise, just take any countable basis on $\mathcal{L}_1(\mathrm{P}, \mathbb{R})$ and project each of its members to the box $[-\gamma, \gamma]$; then, for every $\zeta \in \mathbb{A}_\gamma^i$, there exists a subsequence $\{\lambda_n\}_{n \in \mathcal{K}}, \mathcal{K} \subseteq \mathbb{N}$ converging to ζ in $\mathcal{L}_1$, and in fact

$$0 \leq \|\mathrm{proj}_{[-\gamma,\gamma]}(\lambda_n) - \zeta\|_{\mathcal{L}_1} = \|\mathrm{proj}_{[-\gamma,\gamma]}(\lambda_n) - \mathrm{proj}_{[-\gamma,\gamma]}(\zeta)\|_{\mathcal{L}_1} \leq \|\lambda_n - \zeta\|_{\mathcal{L}_1} \underset{n \to \infty}{\longrightarrow} 0,$$

due to nonexpansiveness of the projection map, and so everything works).

For every $E \in \mathscr{B}(\mathcal{H})$, it follows that

$$\begin{aligned} \|\boldsymbol{G}_{\boldsymbol{p}}(E)\|_{\ell_\infty} &= \sup_{n \in \mathbb{N}} \|\mathbb{E}\{\mathbb{1}_E(\boldsymbol{H}) \lambda_n(\boldsymbol{H}) \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})\}\|_\infty \\ &\leq \sup_{n \in \mathbb{N}} \|\mathbb{E}\{\mathbb{1}_E(\boldsymbol{H}) |\lambda_n(\boldsymbol{H})| |\boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})|\}\|_\infty \\ &\leq \gamma \|\mathbb{E}\{\mathbb{1}_E(\boldsymbol{H}) |\boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})|\}\|_\infty < \infty, \end{aligned}$$

implying that $\boldsymbol{G}_{\boldsymbol{p}}(E)$ is an element of $\mathbb{X}$ for every qualifying E. Evidently, $\boldsymbol{G}_{\boldsymbol{p}}$ is countably additive, and can be represented as a Bochner integral as

$$\boldsymbol{G}_{\boldsymbol{p}}(E) = \int_E \boldsymbol{\Lambda}_{\mathbb{B}}(\boldsymbol{h}) \otimes \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{h}), \boldsymbol{h}) \mathrm{dP}(\boldsymbol{h}), \quad E \in \mathscr{B}(\mathcal{H}),$$

where "$\otimes$" denotes the Kronecker product, and where one can verify that $\boldsymbol{\Lambda}_{\mathbb{B}} \triangleq [\lambda_0\, \lambda_1\, \ldots\, \lambda_n\, \ldots] \in \mathbb{X}$ (see, e.g., [Diestel and Uhl, 1977, Example II.2.10]); note that $\boldsymbol{\Lambda}_{\mathbb{B}}(\boldsymbol{h}) \otimes \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{h}), \boldsymbol{h})$ is indeed in $\mathbb{X}$ and Bochner integrable as well (see proof of Lemma 3 below). We also use the probabilistic notation

$$\boldsymbol{G}_{\boldsymbol{p}}(E) = \mathbb{E}\{\mathbb{1}_E(\boldsymbol{H}) \boldsymbol{\Lambda}_{\mathbb{B}}(\boldsymbol{H}) \otimes \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})\}, \quad E \in \mathscr{B}(\mathcal{H}),$$

with an understanding that expectation here is in the Bochner sense (i.e., expectation of a random element taking values in an infinite-dimensional Banach space).

Using the construction above, and together with another feasible policy $\boldsymbol{p}' \in \Pi$, we define another vector measure $\boldsymbol{G} : \mathscr{B}(\mathcal{H}) \to \mathbb{X}$ as

$$\boldsymbol{G}(E) \triangleq \mathbb{E}\left\{\mathbb{1}_E(\boldsymbol{H})\, \boldsymbol{\Lambda}_{\mathbb{B}}(\boldsymbol{H}) \otimes \begin{bmatrix} \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H}) \\ \boldsymbol{f}(\boldsymbol{p}'(\boldsymbol{H}), \boldsymbol{H}) \end{bmatrix}\right\}, \quad E \in \mathscr{B}(\mathcal{H}),$$

which will serve as our main construction for the rest of the analysis. Observe that $\boldsymbol{G}$ is essentially an interleaved concatenation of the vector measures $\boldsymbol{G}_{\boldsymbol{p}}$ and $\boldsymbol{G}_{\boldsymbol{p}'}$, as defined above. The next key result follows, concerning the structure of the range of $\boldsymbol{G}$.

Lemma 3. *The norm closure of the range of* $\boldsymbol{G}$

$$\boldsymbol{G}(\mathscr{B}(\mathcal{H})) = \{\boldsymbol{x} \in \mathbb{X} | \boldsymbol{x} = \boldsymbol{G}(E),\ \textit{for some } E \in \mathscr{B}(\mathcal{H})\}$$

is convex and norm-compact.

Proof of Lemma 3. We need to verify the conditions under which Theorem 2 (Uhl) is valid. First, it is true that

$$\begin{aligned}
\mathbb{E}\left\{\left\|\boldsymbol{\Lambda}_{\mathbb{B}}(\boldsymbol{H}) \otimes \begin{bmatrix} \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H}) \\ \boldsymbol{f}(\boldsymbol{p}'(\boldsymbol{H}), \boldsymbol{H}) \end{bmatrix}\right\|_{\ell_\infty}\right\} &= \mathbb{E}\left\{\sup_{n \in \mathbb{N}} \left\|\lambda_n(\boldsymbol{H}) \begin{bmatrix} \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H}) \\ \boldsymbol{f}(\boldsymbol{p}'(\boldsymbol{H}), \boldsymbol{H}) \end{bmatrix}\right\|_\infty\right\} \\
&\le \mathbb{E}\left\{\sup_{n \in \mathbb{N}} |\lambda_n(\boldsymbol{H})| \left\|\begin{bmatrix} \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H}) \\ \boldsymbol{f}(\boldsymbol{p}'(\boldsymbol{H}), \boldsymbol{H}) \end{bmatrix}\right\|_\infty\right\} \\
&\le \gamma \mathbb{E}\left\{\left\|\begin{bmatrix} \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H}) \\ \boldsymbol{f}(\boldsymbol{p}'(\boldsymbol{H}), \boldsymbol{H}) \end{bmatrix}\right\|_\infty\right\} \\
&\le \gamma \mathbb{E}\left\{\left\|\begin{bmatrix} \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H}) \\ \boldsymbol{f}(\boldsymbol{p}'(\boldsymbol{H}), \boldsymbol{H}) \end{bmatrix}\right\|_1\right\} < \infty.
\end{aligned}$$

This shows that

$$\boldsymbol{\Lambda}_{\mathbb{B}}(\boldsymbol{H}) \otimes \begin{bmatrix} \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H}) \\ \boldsymbol{f}(\boldsymbol{p}'(\boldsymbol{H}), \boldsymbol{H}) \end{bmatrix} \in \mathcal{L}_1(\mathrm{P}, \mathbb{X}),$$

which also implies that $\boldsymbol{G}$ has a Radon-Nikodym representation. Let us show that $\boldsymbol{G}$ is of bounded variation as well, for essentially the same reasons. According to our definitions, $\boldsymbol{G}$ will be of bounded variation if $|\boldsymbol{G}|(\mathcal{H}) < \infty$. We have

$$|\boldsymbol{G}|(\mathcal{H}) = \sup_{\pi \text{ is a finite partition of } \mathcal{H}} \sum_{A \in \pi} \|\boldsymbol{G}(A)\|_{\ell_\infty}$$

$$
\begin{aligned}
&= \sup_{\pi \text{ is a finite partition of } \mathcal{H}} \sum_{A\in\pi} \left\| \mathbb{E}\left\{ \mathbb{1}_A(\boldsymbol{H})\, \boldsymbol{\Lambda}_{\mathbb{B}}(\boldsymbol{H}) \otimes \begin{bmatrix} \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}),\boldsymbol{H}) \\ \boldsymbol{f}(\boldsymbol{p}'(\boldsymbol{H}),\boldsymbol{H}) \end{bmatrix} \right\} \right\|_{\ell_\infty} \\
&= \sup_{\pi \text{ is a finite partition of } \mathcal{H}} \sum_{A\in\pi} \sup_{n\in\mathbb{N}} \left\| \mathbb{E}\left\{ \mathbb{1}_A(\boldsymbol{H})\, \lambda_n(\boldsymbol{H}) \begin{bmatrix} \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}),\boldsymbol{H}) \\ \boldsymbol{f}(\boldsymbol{p}'(\boldsymbol{H}),\boldsymbol{H}) \end{bmatrix} \right\} \right\|_{\infty} \\
&\le \sup_{\pi \text{ is a finite partition of } \mathcal{H}} \gamma \sum_{A\in\pi} \mathbb{E}\left\{ \mathbb{1}_A(\boldsymbol{H}) \left\| \begin{bmatrix} \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}),\boldsymbol{H}) \\ \boldsymbol{f}(\boldsymbol{p}'(\boldsymbol{H}),\boldsymbol{H}) \end{bmatrix} \right\|_{\infty} \right\} \\
&= \gamma \mathbb{E}\left\{ \left\| \begin{bmatrix} \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}),\boldsymbol{H}) \\ \boldsymbol{f}(\boldsymbol{p}'(\boldsymbol{H}),\boldsymbol{H}) \end{bmatrix} \right\|_{\infty} \right\} < \infty.
\end{aligned}
$$

This shows that the vector measure $\boldsymbol{G}$ is of bounded variation. To show that $\boldsymbol{G}$ is nonatomic, consider its primitive construction and suppose that $E \in \mathscr{B}(\mathcal{H})$ is such that

$$
\boldsymbol{G}(E) \neq \boldsymbol{0} \iff \mathbb{E}\{\mathbb{1}_E(\boldsymbol{H})\lambda_n(\boldsymbol{H}) f_i(\widetilde{\boldsymbol{p}}(\boldsymbol{H}),\boldsymbol{H})\} \neq 0, \text{ for some } n \in \mathbb{N}, i \in \mathbb{N}_N^+ \text{ and for } \widetilde{\boldsymbol{p}} = \boldsymbol{p} \text{ or } \boldsymbol{p}'.
$$

Note that we necessarily have $\mathrm{P}(E) > 0$ (for otherwise $\boldsymbol{G}(E) = \boldsymbol{0}$). Without loss of generality take $\widetilde{\boldsymbol{p}} = \boldsymbol{p}$ and $n = i = 1$. Consider the restriction $(E, \mathscr{B}(E))$ of $(\mathcal{H}, \mathscr{B}(\mathcal{H}))$ on E. Then, by a lemma of Blackwell [Blackwell, 1951, Lemma], nonatomicity of P implies the existence of a Borel subset $E' \subseteq E$ such that $\mathrm{P}(E) > \mathrm{P}(E') = \mathrm{P}(E)/2 > 0$, for which

$$
\mathbb{E}\{\mathbb{1}_{E'}(\boldsymbol{H})\lambda_1(\boldsymbol{H}) f_1(\boldsymbol{p}(\boldsymbol{H}),\boldsymbol{H})\} = \frac{1}{2}\mathbb{E}\{\mathbb{1}_{E}(\boldsymbol{H})\lambda_1(\boldsymbol{H}) f_1(\boldsymbol{p}(\boldsymbol{H}),\boldsymbol{H})\} \neq 0.
$$

This necessarily implies that $\boldsymbol{G}(E') \neq \boldsymbol{0}$ as well as $\boldsymbol{G}(E \backslash E') \neq \boldsymbol{0}$. By definition, it follows that $\boldsymbol{G}$ has no atoms. Consequently, the conditions of Theorem 2 are fulfilled. Enough said. ■

The conclusions of Lemma 3 are sufficient to ensure the existence of a feasible policy $\boldsymbol{p}_\alpha \in \Pi$, such that (5) is true. To see this, let us first consider the range $\boldsymbol{G}(\mathscr{B}(\mathcal{H}))$ of $\boldsymbol{G}$. Of course $\boldsymbol{x} = \boldsymbol{G}(\mathcal{H})$ and $\boldsymbol{x}' = \boldsymbol{G}(\emptyset) = \boldsymbol{0}$ are both elements of $\boldsymbol{G}(\mathscr{B}(\mathcal{H}))$. Therefore, Lemma 3 implies that, for every $\alpha \in [0,1]$, the convex combination $\alpha\boldsymbol{x} + (1-\alpha)\boldsymbol{x}' \equiv \alpha\boldsymbol{x}$ lies in the norm closure of $\boldsymbol{G}(\mathscr{B}(\mathcal{H}))$. In order words, for each α, there exists a sequence of events $\{E_n^\alpha \in \mathscr{B}(\mathcal{H})\}_{n\in\mathbb{N}}$ such that

$$
\lim_{n\to\infty} \|\alpha\boldsymbol{x} - \boldsymbol{G}(E_n^\alpha)\|_{\ell_\infty} = 0.
$$

This in particular implies that (note that limits here are with respect to the natural norm of $\mathbb{X}$)

$$
\lim_{n\to\infty} \|\alpha\boldsymbol{G}_{\boldsymbol{p}}(\mathcal{H}) - \boldsymbol{G}_{\boldsymbol{p}}(E_n^\alpha)\|_{\ell_\infty} = 0,
$$

and by a symmetric argument,

$$
\lim_{n\to\infty} \|(1-\alpha)\boldsymbol{G}_{\boldsymbol{p}'}(\mathcal{H}) - \boldsymbol{G}_{\boldsymbol{p}'}((E_n^\alpha)^c)\|_{\ell_\infty} = 0.
$$

Now, we may define the sequence of policies

$$
\boldsymbol{p}_\alpha^n(\boldsymbol{H}) = \mathbb{1}_{E_n^\alpha}(\boldsymbol{H})\boldsymbol{p}(\boldsymbol{H}) + \mathbb{1}_{\mathcal{H}\backslash E_n^\alpha}(\boldsymbol{H})\boldsymbol{p}'(\boldsymbol{H}) = \begin{cases} \boldsymbol{p}(\boldsymbol{H}), & \text{if } \boldsymbol{H} \in E_n^\alpha \\ \boldsymbol{p}'(\boldsymbol{H}), & \text{if } \boldsymbol{H} \in \mathcal{H}\backslash E_n^\alpha \end{cases}, \quad n \in \mathbb{N}.
$$

Of course, it holds that $\boldsymbol{p}_\alpha^n(\boldsymbol{H}) \in \Pi$ for all $n \in \mathbb{N}$ because Π is decomposable. Then, it follows that

$$
\|\boldsymbol{G}_{\boldsymbol{p}_\alpha^n}(\mathcal{H}) - \alpha\boldsymbol{G}_{\boldsymbol{p}}(\mathcal{H}) - (1-\alpha)\boldsymbol{G}_{\boldsymbol{p}'}(\mathcal{H})\|_{\ell_\infty}
$$

$$
\begin{aligned}
&= \|\boldsymbol{G}_{\boldsymbol{p}_\alpha^n}(E_n^\alpha) + \boldsymbol{G}_{\boldsymbol{p}_\alpha^n}((E_n^\alpha)^c) - \alpha \boldsymbol{G}_{\boldsymbol{p}}(\mathcal{H}) - (1-\alpha)\boldsymbol{G}_{\boldsymbol{p}'}(\mathcal{H})\|_{\ell_\infty} \\
&= \|\boldsymbol{G}_{\boldsymbol{p}}(E_n^\alpha) + \boldsymbol{G}_{\boldsymbol{p}'}((E_n^\alpha)^c) - \alpha \boldsymbol{G}_{\boldsymbol{p}}(\mathcal{H}) - (1-\alpha)\boldsymbol{G}_{\boldsymbol{p}'}(\mathcal{H})\|_{\ell_\infty} \\
&\le \|\boldsymbol{G}_{\boldsymbol{p}}(E_n^\alpha) - \alpha \boldsymbol{G}_{\boldsymbol{p}}(\mathcal{H})\|_{\ell_\infty} + \|\boldsymbol{G}_{\boldsymbol{p}'}((E_n^\alpha)^c) - (1-\alpha)\boldsymbol{G}_{\boldsymbol{p}'}(\mathcal{H})\|_{\ell_\infty},
\end{aligned}
$$

which implies that

$$
\lim_{n\to\infty} \|\boldsymbol{G}_{\boldsymbol{p}_\alpha^n}(\mathcal{H}) - \alpha \boldsymbol{G}_{\boldsymbol{p}}(\mathcal{H}) - (1-\alpha)\boldsymbol{G}_{\boldsymbol{p}'}(\mathcal{H})\|_{\ell_\infty} = 0.
$$

Equivalently, we have shown that for every $\varepsilon > 0$, there exists a positive number $N(\varepsilon) > 0$, such that for every $n > N(\varepsilon)$,

$$
\begin{aligned}
&\|\mathbb{E}\{\boldsymbol{\Lambda}_{\mathbb{B}}(\boldsymbol{H}) \otimes \boldsymbol{f}(\boldsymbol{p}_\alpha^n(\boldsymbol{H}), \boldsymbol{H})\} \\
&\quad - \alpha \mathbb{E}\{\boldsymbol{\Lambda}_{\mathbb{B}}(\boldsymbol{H}) \otimes \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})\} - (1-\alpha)\mathbb{E}\{\boldsymbol{\Lambda}_{\mathbb{B}}(\boldsymbol{H}) \otimes \boldsymbol{f}(\boldsymbol{p}'(\boldsymbol{H}), \boldsymbol{H})\}\|_{\ell_\infty} \le \varepsilon.
\end{aligned}
$$

Evidently, $N(\varepsilon)$ is uniform over the individual elements of the countable basis $\mathbb{B}$. We can rewrite the preceding expression as

$$
\begin{aligned}
&|\mathbb{E}\{\lambda_m(\boldsymbol{H}) f_i(\boldsymbol{p}_\alpha^n(\boldsymbol{H}), \boldsymbol{H})\} \\
&\quad - \alpha \mathbb{E}\{\lambda_m(\boldsymbol{H}) f_i(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})\} - (1-\alpha)\mathbb{E}\{\lambda_m(\boldsymbol{H}) f_i(\boldsymbol{p}'(\boldsymbol{H}), \boldsymbol{H})\}| \le \varepsilon,
\end{aligned}
$$

for every pair $(m, i) \in \mathbb{N} \times \mathbb{N}_N^+$. Now, for each choice of $\zeta \in \mathbb{A}_\gamma^i$, $i \in \mathbb{N}_N^+$, we can extract a subsequence $\{\lambda_m\}_{m\in\mathcal{K}}, \mathcal{K} \subseteq \mathbb{N}$ converging to ζ in $\mathcal{L}_1$. A consequence of this is the existence of a further sub-subsequence $\{\lambda_m\}_{m\in\mathcal{K}'}$, $\mathcal{K}' \subseteq \mathcal{K}$ such that

$$
\lambda_m \underset{\mathcal{K}' \ni m \to \infty}{\longrightarrow} \zeta, \quad \text{P-a.e.}
$$

Then, by dominated convergence, we have for each i-th element of the service vector $\boldsymbol{f}$,

$$
\begin{aligned}
\mathbb{E}\{\lambda_m(\boldsymbol{H}) f_i(\boldsymbol{p}_\alpha^n(\boldsymbol{H}), \boldsymbol{H})\} &\underset{\mathcal{K}' \ni m \to \infty}{\longrightarrow} \mathbb{E}\{\zeta(\boldsymbol{H}) f_i(\boldsymbol{p}_\alpha^n(\boldsymbol{H}), \boldsymbol{H})\}, \\
\mathbb{E}\{\lambda_m(\boldsymbol{H}) f_i(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})\} &\underset{\mathcal{K}' \ni m \to \infty}{\longrightarrow} \mathbb{E}\{\zeta(\boldsymbol{H}) f_i(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})\} \quad \text{and} \\
\mathbb{E}\{\lambda_m(\boldsymbol{H}) f_i(\boldsymbol{p}'(\boldsymbol{H}), \boldsymbol{H})\} &\underset{\mathcal{K}' \ni m \to \infty}{\longrightarrow} \mathbb{E}\{\zeta(\boldsymbol{H}) f_i(\boldsymbol{p}'(\boldsymbol{H}), \boldsymbol{H})\}.
\end{aligned}
$$

Therefore, we have that, for every $\varepsilon > 0$ there exists a positive number $N(\varepsilon) > 0$, such that for every $n > N(\varepsilon)$ and for every $\zeta \in \mathbb{A}_\gamma^i$,

$$
\begin{aligned}
&|\mathbb{E}\{\zeta(\boldsymbol{H}) f_i(\boldsymbol{p}_\alpha^n(\boldsymbol{H}), \boldsymbol{H})\} \\
&\quad - \alpha \mathbb{E}\{\zeta(\boldsymbol{H}) f_i(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})\} - (1-\alpha)\mathbb{E}\{\zeta(\boldsymbol{H}) f_i(\boldsymbol{p}'(\boldsymbol{H}), \boldsymbol{H})\}| \le \varepsilon, \quad i \in \mathbb{N}_N^+.
\end{aligned}
$$

The last expression implies in particular that

$$
\begin{aligned}
&\Bigg| \inf_{\zeta \in \mathbb{A}_\gamma^i} \mathbb{E}\{\zeta(\boldsymbol{H}) f_i(\boldsymbol{p}_\alpha^n(\boldsymbol{H}), \boldsymbol{H})\} \\
&\quad - \inf_{\zeta \in \mathbb{A}_\gamma^i} \mathbb{E}\{\zeta(\boldsymbol{H})[\alpha f_i(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H}) + (1-\alpha) f_i(\boldsymbol{p}'(\boldsymbol{H}), \boldsymbol{H})]\} \Bigg| \le \varepsilon, \quad i \in \mathbb{N}_N^+,
\end{aligned}
$$

which is the same as

$$
\left\| \inf_{\boldsymbol{\zeta} \in \mathbb{A}_\gamma^S} \mathbb{E}\{\boldsymbol{\zeta}(\boldsymbol{H}) \circ \boldsymbol{f}(\boldsymbol{p}_\alpha^n(\boldsymbol{H}), \boldsymbol{H})\} - \inf_{\boldsymbol{\zeta} \in \mathbb{A}_\gamma^S} \mathbb{E}\{\boldsymbol{\zeta}(\boldsymbol{H}) \circ [\alpha \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H}) + (1-\alpha) \boldsymbol{f}(\boldsymbol{p}'(\boldsymbol{H}), \boldsymbol{H})]\} \right\|_\infty \le \varepsilon.
$$

By risk duality, we obtain that, for every $\varepsilon > 0$, there is $N(\varepsilon) > 0$ such that for every $n > N(\varepsilon)$, it is true that

$$\| -\boldsymbol{\rho}(-\boldsymbol{f}(\boldsymbol{p}_\alpha^n(\boldsymbol{H}),\boldsymbol{H})) + \boldsymbol{\rho}(-\alpha \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}),\boldsymbol{H}) - (1-\alpha)\boldsymbol{f}(\boldsymbol{p}'(\boldsymbol{H}),\boldsymbol{H}))\|_\infty \leq \varepsilon.$$

This verifies (5) for every choice of $\boldsymbol{p} \in \Pi$, $\boldsymbol{p}' \in \Pi$ and for every $\alpha \in [0,1]$, and thus proves that the set $\mathcal{C}$ is convex. Indeed, since by assumption the risk limit span

$$\mathbb{S} \triangleq \left\{\boldsymbol{s} \in \mathbb{R}^N \middle| \boldsymbol{s} = \boldsymbol{\rho}(-\boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}),\boldsymbol{H})),\ \text{for some } \boldsymbol{p} \in \Pi\right\}$$

is a closed set, then the unique limit point of the convergent sequence $\{\boldsymbol{\rho}(-\boldsymbol{f}(\boldsymbol{p}_\alpha^n(\boldsymbol{H}),\boldsymbol{H}))\}_{n\in\mathbb{N}}$ must be attained at $\boldsymbol{\rho}(-\boldsymbol{f}(\boldsymbol{p}_\alpha(\boldsymbol{H}),\boldsymbol{H}))$ for some $\boldsymbol{p}_\alpha \in \Pi$, and cannot be other than

$$\boldsymbol{\rho}(-\boldsymbol{f}(\boldsymbol{p}_\alpha(\boldsymbol{H}),\boldsymbol{H})) \equiv \boldsymbol{\rho}(-\alpha \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}),\boldsymbol{H}) - (1-\alpha)\boldsymbol{f}(\boldsymbol{p}'(\boldsymbol{H}),\boldsymbol{H})).$$

Therefore, by convexity and positive homogeneity of $\boldsymbol{\rho}$, we have

$$\begin{aligned}-\boldsymbol{\rho}(-\boldsymbol{f}(\boldsymbol{p}_\alpha^n(\boldsymbol{H}),\boldsymbol{H})) &\geq -\boldsymbol{\rho}(-\alpha \boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}),\boldsymbol{H})) - \boldsymbol{\rho}(-(1-\alpha)\boldsymbol{f}(\boldsymbol{p}'(\boldsymbol{H}),\boldsymbol{H})) \\ &= -\alpha\boldsymbol{\rho}(-\boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}),\boldsymbol{H})) - (1-\alpha)\boldsymbol{\rho}(-\boldsymbol{f}(\boldsymbol{p}'(\boldsymbol{H}),\boldsymbol{H})).\end{aligned}$$

The proof is now complete. ■

Remark 1. Note that the elementary fact that enables interchanging the order of limiting operations relative to n and m above is that convergence over n is uniform over m, i.e., the index of the elements in the countable basis $\mathbb{B}$. Then, we reiterate the same procedure for every member ζ of each of the risk envelopes $\mathbb{A}_\gamma^i$, $i \in \mathbb{N}_N^+$, by extracting a different subsequence out of $\mathbb{B}$ each time.

4.3 Convexity of $\mathcal{C}$ Implies Strong Duality

Let us now finish the proof of Theorem 1 by exploiting the convexity of the utility-constraint set

$$\mathcal{C} = \left\{(\delta_o, \boldsymbol{\delta}_r, \boldsymbol{\delta}_d) \middle| \begin{array}{l} g^o(\boldsymbol{x}) \geq \delta_o \\ -\boldsymbol{\rho}(-\boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}),\boldsymbol{H})) - \boldsymbol{x} \geq \boldsymbol{\delta}_r, \\ \boldsymbol{g}(\boldsymbol{x}) \geq \boldsymbol{\delta}_d \end{array} \text{for some } (\boldsymbol{x},\boldsymbol{p}) \in \mathcal{X} \times \Pi \right\},$$

the expression of which we repeat here for convenience, together with condition **(5)** of Assumption 1, namely that problem (RCP) satisfies Slater's condition. Let us recall the Lagrangian associated with problem (RCP), i.e.,

$$\mathsf{L}(\boldsymbol{x},\boldsymbol{p},\boldsymbol{\lambda}) = g^o(\boldsymbol{x}) + \langle \boldsymbol{\lambda}_{\boldsymbol{g}}, \boldsymbol{g}(\boldsymbol{x})\rangle + \langle \boldsymbol{\lambda}_{\boldsymbol{\rho}}, -\boldsymbol{\rho}(-\boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}),\boldsymbol{H})) - \boldsymbol{x}\rangle$$

where $(\boldsymbol{x},\boldsymbol{p},\boldsymbol{\lambda}) \in \mathbf{C} \times \Pi \times \mathbb{R}^{N_{\boldsymbol{g}}} \times \mathbf{C}$, for which we already know that

$$\inf_{\boldsymbol{\lambda}\geq \mathbf{0}} \sup_{(\boldsymbol{x},\boldsymbol{p})\in\mathcal{X}\times\Pi} \mathsf{L}(\boldsymbol{x},\boldsymbol{p},\boldsymbol{\lambda}) = \mathsf{D}^* \geq \mathsf{P}^*.$$

Strong duality of (RCP) will be immediate if we can show that $\mathsf{D}^* \leq \mathsf{P}^*$.

Our discussion closely follows [Chamon et al., 2020, Theorem 1 and its proof] and [Chamon et al., 2021, Appendix A], and is based on a simple and standard application of the supporting hyperplane theorem (see e.g., [Bertsekas, 2009, Proposition 1.5.1]), which we outline below for completeness. Note that the same technique would be applicable in case our initial problem (RCP) was originally convex. The reason is that the only fact needed at this point is the "top-level" convexity of the set $\mathcal{C}$, which would come for free if (RCP) were itself a convex program.

Theorem 3 (**Supporting Hyperplane Theorem**). *Let $\mathcal{A} \subseteq \mathbb{R}^n$ be a nonempty convex set. If $\boldsymbol{\delta}^* \in \mathbb{R}^n$ is not in the interior of $\mathcal{A}$, then there exists a hyperplane passing through $\boldsymbol{\delta}^*$ such that $\mathcal{A}$ is in one of its closed halfspaces. In other words, there exists a vector $\boldsymbol{\lambda} \neq \mathbf{0}$ such that, for every $\boldsymbol{\delta} \in \mathcal{A}$, it holds that $\langle \boldsymbol{\lambda}, \boldsymbol{\delta}^* \rangle \geq \langle \boldsymbol{\lambda}, \boldsymbol{\delta} \rangle$.*

Remark 2. Note that the inequality in the supporting hyperplane theorem can be equivalently reversed by flipping the sign of the support vector $\boldsymbol{\lambda}$.

We start by observing that since (RCP) satisfies Slater's condition, it follows that $\mathcal{C}$ is nonempty. So $\mathcal{C}$ is a nonempty convex set. Also observe that the point $(\mathsf{P}^*, \mathbf{0}, \mathbf{0})$ cannot be in the interior of $\mathcal{C}$, for otherwise there would exist an $\varepsilon > 0$ such that $(\mathsf{P}^* + \varepsilon, \mathbf{0}, \mathbf{0})$, leading to a contradiction. We can then apply the supporting hyperplane theorem on the pair $\mathcal{C}$ and $(\mathsf{P}^*, \mathbf{0}, \mathbf{0})$, implying existence of a vector of multipliers $(\lambda_o, \boldsymbol{\lambda_g}, \boldsymbol{\lambda_\rho}) \neq \mathbf{0}$ such that, for every $(\delta_o, \boldsymbol{\delta}_r, \boldsymbol{\delta}_d) \in \mathcal{C}$, it is true that

$$\lambda_o \delta_o + \langle \boldsymbol{\lambda_g}, \boldsymbol{\delta}_r \rangle + \langle \boldsymbol{\lambda_\rho}, \boldsymbol{\delta}_d \rangle \leq \lambda_o \mathsf{P}^*.$$

It readily follows that, in fact, $(\lambda_o, \boldsymbol{\lambda_g}, \boldsymbol{\lambda_\rho}) \geq \mathbf{0}$. Indeed, if any component of $(\lambda_o, \boldsymbol{\lambda_g}, \boldsymbol{\lambda_\rho})$ was negative, then we could choose $(\delta_o, \boldsymbol{\delta}_r, \boldsymbol{\delta}_d) \in \mathcal{C}$ such that the corresponding inner product becomes arbitrarily large (note that $\mathcal{C}$ is unbounded below), eventually violating the inequality above, regardless of the sign of P^*.

The second fact we may show is that $\lambda_o \neq 0$, which implies that $\lambda_o > 0$. Again, if $\lambda_o = 0$, then we would have that

$$\langle \boldsymbol{\lambda_g}, \boldsymbol{\delta}_r \rangle + \langle \boldsymbol{\lambda_\rho}, \boldsymbol{\delta}_d \rangle \leq 0.$$

But $(\boldsymbol{\lambda_g}, \boldsymbol{\lambda_\rho}) \neq \mathbf{0}$ (i.e., there is at least one nonzero component) and $(\boldsymbol{\lambda_g}, \boldsymbol{\lambda_\rho}) \geq \mathbf{0}$ (shown above), and problem (RCP) satisfies Slater's condition, so the preceding inequality is also absurd.

As a result, we may divide by $\lambda_o > 0$, showing that there is $(\boldsymbol{\lambda}^\star_{\boldsymbol{g}} \triangleq \boldsymbol{\lambda_g}/\lambda_o, \boldsymbol{\lambda}^\star_{\boldsymbol{\rho}} \triangleq \boldsymbol{\lambda_\rho}/\lambda_o) \geq \mathbf{0}$, such that, for every $(\delta_o, \boldsymbol{\delta}_r, \boldsymbol{\delta}_d) \in \mathcal{C}$, it holds that

$$\delta_o + \langle \boldsymbol{\lambda}^\star_{\boldsymbol{g}}, \boldsymbol{\delta}_r \rangle + \langle \boldsymbol{\lambda}^\star_{\boldsymbol{\rho}}, \boldsymbol{\delta}_d \rangle \leq \mathsf{P}^*.$$

By construction of the set $\mathcal{C}$, we obtain

$$\mathsf{L}(\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\lambda}^\star) = g^o(\boldsymbol{x}) + \langle \boldsymbol{\lambda}^\star_{\boldsymbol{g}}, \boldsymbol{g}(\boldsymbol{x}) \rangle + \left\langle \boldsymbol{\lambda}^\star_{\boldsymbol{\rho}}, -\boldsymbol{\rho}(-\boldsymbol{f}(\boldsymbol{p}(\boldsymbol{H}), \boldsymbol{H})) - \boldsymbol{x} \right\rangle \leq \mathsf{P}^*,$$

for every feasible pair $(\boldsymbol{x}, \boldsymbol{p}) \in \mathcal{X} \times \Pi$, since all such pairs are included in $\mathcal{C}$. This further implies that we are allowed to maximize both sides of the inequality over all feasible $(\boldsymbol{x}, \boldsymbol{p})$, yielding

$$\begin{aligned} -\infty < \mathsf{D}^* &= \inf_{\boldsymbol{\lambda} \geq \mathbf{0}} \sup_{(\boldsymbol{x}, \boldsymbol{p}) \in \mathcal{X} \times \Pi} \mathsf{L}(\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\lambda}) \\ &\leq \sup_{(\boldsymbol{x}, \boldsymbol{p}) \in \mathcal{X} \times \Pi} \mathsf{L}(\boldsymbol{x}, \boldsymbol{p}, \boldsymbol{\lambda}^\star) \leq \mathsf{P}^*, \end{aligned}$$

and we are done. ■

Acknowledgement

The author would like to thank Dr. Spyridon Pougakiotis for his careful reading and useful suggestions on several earlier drafts of this work.

References

Dimitri P. Bertsekas. *Convex optimization theory.* Athena Scientific, 1st edition, 2009. ISBN 1886529310. URL http://www.athenasc.com/convexduality.html.

David Blackwell. The Range of Certain Vector Integrals. *Proceedings of the American Mathematical Society*, 2(3):390, jun 1951. ISSN 00029939. doi: 10.2307/2031763.

Stephen P. Boyd and Lieven. Vandenberghe. *Convex Optimization.* Cambridge University Press, 2004. ISBN 9780521833783.

Luiz F. O. Chamon, Santiago Paternain, Miguel Calvo-Fullana, and Alejandro Ribeiro. Constrained Learning with Non-Convex Losses. *arXiv preprint, arXiv:2103.05134*, mar 2021. URL http://arxiv.org/abs/2103.05134.

Luiz F.O. Chamon, Yonina C. Eldar, and Alejandro Ribeiro. Functional Nonlinear Sparse Models. *IEEE Transactions on Signal Processing*, 68:2449–2463, 2020. ISSN 19410476. doi: 10.1109/TSP.2020.2982834.

J. Diestel and J. R. Uhl. *Vector Measures.* American Mathematical Society, Providence, math. surv edition, 1977.

Zhi Quan Luo and Shuzhong Zhang. Dynamic Spectrum Management: Complexity and Duality. *IEEE Journal on Selected Topics in Signal Processing*, 2(1):57–73, feb 2008. ISSN 19324553. doi: 10.1109/JSTSP.2007.914876.

Włodzimierz Ogryczak and Andrzej Ruszczyński. From Stochastic Dominance to Mean-Risk Models: Semideviations as Risk Measures. *European Journal of Operational Research*, 116(1):33–50, jul 1999. ISSN 03772217. doi: 10.1016/S0377-2217(98)00167-2.

Włodzimierz Ogryczak and Andrzej Ruszczyński. Dual Stochastic Dominance and Related Mean-Risk Models. *SIAM Journal on Optimization*, 13(1):60–78, 2002. ISSN 1052-6234. doi: 10.1137/S1052623400375075.

Alejandro Ribeiro. Optimal Resource Allocation in Wireless Communication and Networking. *EURASIP Journal on Wireless Communications and Networking*, 2012(1):272, 2012. doi: 10.1186/1687-1499-2012-272. URL http://jwcn.eurasipjournals.com/content/2012/1/272.

R. Tyrrell Rockafellar and Stanislav Uryasev. Optimization of Conditional Value-at-Risk. *Journal of Risk*, 2:21–41, 2000. ISSN 10769986. doi: 10.2307/1165345.

Andrzej P. Ruszczyński. *Nonlinear Optimization.* Princeton University Press, 2006. ISBN 9780691119151. URL https://press.princeton.edu/titles/8219.html.

Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczyński. *Lectures on Stochastic Programming: Modeling and Theory.* Society for Industrial and Applied Mathematics, 2nd edition, 2014. ISBN 089871687X. doi: http://dx.doi.org/10.1137/1.9780898718751.

J. J. Uhl. The Range of a Vector-Valued Measure. *Proceedings of the American Mathematical Society*, 23(1):158, oct 1969. ISSN 00029939. doi: 10.2307/2037509.