Note: Designing Flexible Systems Using a New Notion of Submodularity

Achal Bassamboo
Kellogg School of Management, Northwestern University, a-bassamboo@northwestern.edu

Leon Yang Chu
Marshall School of Business, University of Southern California, leonyzhu@usc.edu

Ramandeep S. Randhawa
Marshall School of Business, University of Southern California, ramandeep.randhawa@marshall.usc.edu

September 7, 2011

We study the problem of optimal flexibility capacity portfolio selection by introducing a new notion of submodularity in correspondences, which extends the classical notion of submodular functions. In particular, we prove that the correspondence that maps flexible resources to the set of demands that they can process is submodular, and use the properties of submodular correspondences to compare different flexibility configurations and derive insights into the optimal capacity portfolio.

Key words: Flexibility, submodularity, correspondences

1. Introduction

We study the problem of capacity selection for a firm that produces n different types of products and has access to resources with different production capabilities. The ability of a resource to produce different product types is termed as its flexibility. Jordan and Graves (1995) studied this problem using a newsvendor framework in which they proposed a chaining configuration consisting of bi-flexible resources (that can produce two product types), and showed that chaining configurations can perform quite well.

In this paper, we study this classical problem with the goal of developing a means of comparing different resource configurations that can be used to draw insights into the structure of optimal resource configurations. To do so, we first extend the notion of submodular functions from real-valued functions to correspondences, and then use properties of these correspondences to compare and rank the performance of different resource configurations. An important feature of this ranking is that it is *independent* of the manner in which the cost of a demand shortfall is computed. Under some conditions on the costs of investing in resources (for instance, linear economies of scope), this ranking yields a nice structure on the optimal capacity portfolio: resources invested in must have "similar" flexibility, and resources with lower flexibility must be more productive than those with higher flexibility (see Proposition 5 in Section 3.4 for a formal characterization).

This paper consists of two main sections: Section 2 introduces the notion of submodular correspondences, and Section 3 uses this notion in optimizing flexible capacity portfolios. The focus of this paper is on single-period newsvendor models and we discuss the implications of our results for other models in Section 4.

Notation. We use $\mathcal{P}(A)$ to denote the power set of A. For sets A and B, we use the notation $A+B \equiv \{x+y|x \in A, y \in B\}$. Further, we use the notation $m \times A = A + (m-1) \times A$ for $m \in Z, m > 1$ and $mA \equiv \{mx|x \in A\}$ for $m \in \mathbb{R}_+$. For vectors $x, y \in \mathbb{R}^n$, we use $x \vee y$ to denote the component-wise maximum and $x \wedge y$ to denote the component-wise minimum of x and y. Finally, we say that a set $X \subseteq \mathbb{R}^n$ is a lattice, if for any $x, y \in X$ we have $x \vee y, x \wedge y \in X$.

2. Submodularity for Correspondences

We begin by recalling the definition of submodularity for real-valued functions. Let $X \subseteq \mathbb{R}^n$ be a lattice. Then, a function $f: X \to \mathbb{R}$ is submodular if

$$f(x \vee y) + f(x \wedge y) \leq f(x) + f(y), \text{ for all } x, y \in X.$$

Further, if $-f$ is submodular then f is said to be supermodular, and if the inequality is changed to an equality the function is said to be modular. For a more detailed discussion on submodularity, see Chapter 2 of Simchi-Levi, Chen and Bramel (2005).

We extend this notion of submodularity to correspondences as follows:

Definition 1. A correspondence $\mathcal{F}: X \to \mathcal{P}(\mathbb{R}^n)$ is *submodular* if

$$\mathcal{F}(x \vee y) + \mathcal{F}(x \wedge y) \subseteq \mathcal{F}(x) + \mathcal{F}(y). \tag{1}$$

If the set inclusion is reversed in (1), then the correspondence $\mathcal{F}$ is said to be supermodular, and if both sets are equal then $\mathcal{F}$ is said to be modular. This new notion is a generalization because a function $f: X \to \mathbb{R}$ is submodular if and only if the correspondence $\mathcal{F}(x) = (-\infty, f(x))$ is submodular. The following are some properties of submodular correspondences.

Proposition 1 **(Properties of submodular correspondences)**.

1. *If a correspondence $\mathcal{F}_i: X \to \mathcal{P}(\mathbb{R}^n)$ is submodular, then $\mathcal{F} = \sum_{i=1}^n \alpha_i \mathcal{F}_i$ is submodular for $\alpha_i \in \mathbb{R}_+$ for $i = 1, 2, \ldots, n$.*

2. *If a correspondence $\mathcal{F}: X \to \mathcal{P}(\mathbb{R}^n)$ is submodular, then the convex hull of $\mathcal{F}$ is also submodular.*

2.1. Illustration 1: A submodular correspondence.

Define $X = \{(0,0),(1,0),(0,1),(1,1)\}$ and the correspondence $\mathcal{F}(\cdot)$ on X as follows

$$\mathcal{F}(x) = \begin{cases} \{(0,0)\}, & \text{if } x = (0,0), \\ \{(0,0),(1,0)\}, & \text{if } x = (1,0), \\ \{(0,0),(0,1)\}, & \text{if } x = (0,1), \\ \{(0,0),(1,0),(0,1)\}, & \text{if } x = (1,1). \end{cases}$$

Then, $\mathcal{F}$ is a submodular correspondence. To see this, observe that for all $x, y \in X$ such that $\{x,y\} \neq \{(0,1),(1,0)\}$, we either have $x \vee y = x$ and $x \wedge y = y$ or $x \vee y = y$ and $x \wedge y = x$, which implies that $\mathcal{F}(x \vee y) + \mathcal{F}(x \wedge y) = \mathcal{F}(x) + \mathcal{F}(y)$. Thus, to ascertain the submodularity of $\mathcal{F}$, we only need to check relation (1) for $x = (0,1)$ and $y = (1,0)$. At these values, we have

$$\mathcal{F}(x) + \mathcal{F}(y) = X \supset \{(0,0),(0,1),(1,0)\} = \mathcal{F}(x \vee y) + \mathcal{F}(x \wedge y).$$

Thus, $\mathcal{F}$ is a submodular correspondence.

3. Submodular Correspondences and Optimal Flexibility Configurations

This section explores the implications of submodular correspondences for selecting optimal flexibility portfolios. Section 3.1 relates the concept of submodular correspondences to the performance of flexible resources. Sections 3.2 and 3.3 then formalize this connection in multi-product newsvendor settings and derive properties of the optimal flexibility portfolio. Finally, Section 3.4 generalizes our results to a setting in which resources can also differ in their productivity or processing efficiency.

3.1. A two-product resource allocation problem in the setting of Illustration 1

The simple submodular correspondence described in Illustration 1, $\mathcal{F}$, (cf. Section 2.1) can be mapped to a classical operations problem of matching supply and demand in a two product setting with three resources: one dedicated resource that can satisfy product-1 demand, one dedicated resource that can satisfy product-2 demand, and one flexible resource that can satisfy either product-1 or product-2 demand, represented by $(1,0)$, $(0,1)$, and $(1,1)$ respectively. (The element (0,0) represents a null resource that satisfies no demand.) The correspondence $\mathcal{F}(r)$ thus enumerates all possible demand vectors that can be met by a resource r, and in this sense represents its *performance region*. For instance, $\mathcal{F}(1,1) = \{(0,0),(0,1),(1,0)\}$ means that the resource $(1,1)$ can either satisfy zero demand, denoted by $(0,0)$, or product 1 demand $(1,0)$, or product 2 demand $(0,1)$. The performance region of a portfolio of resources can then be obtained by simply adding the performance regions for individual resources in the portfolio. That is, for a portfolio consisting of resources $r_1, r_2, \ldots, r_k$, the performance region of the portfolio equals $\sum_{i=1}^{k} \mathcal{F}(r_i)$.

In this two-product setting, it is obvious that, disregarding resource cost considerations, two dedicated resources should be preferred to the flexible resource alone in the sense that the former can satisfy more demand. This relation is mathematically equivalent to the relation $\mathcal{F}(1,0)+\mathcal{F}(0,1)\supset\mathcal{F}(1,1)=\mathcal{F}(1,1)+\mathcal{F}(0,0)$, which is precisely the condition for the submodularity of the performance region $\mathcal{F}$. Thus, in this simple resource allocation scenario, we find that the underlying performance region is a submodular correspondence. We next prove that this property holds for general, multi-product settings as well.

3.2. The model

Consider a system with n product types. There are 2^n-1 different types of potential resources. We represent a resource by the demand types it can process, so a resource r, where $r\in\{0,1\}^n$ processes demand type i if any only if $r_i=1$. Thus, a resource r is more flexible than resource s if $r\geq s$. Further, the number of different products that resource r can process is given by $\sum_{i=1}^n r_i$.

Let $\mathcal{F}:\{0,1\}^n\to\mathcal{P}(\{0,1\}^n)$ denote the correspondence that maps the resources to the demand vectors that they can satisfy. Then, we can write

$$\mathcal{F}(r)=\{y\in\{0,1\}^n|\sum_{i=1}^n y_i\leq 1 \text{ and } y_i\leq r_i \text{ for all } 1\leq i\leq n\}. \tag{2}$$

We then obtain the following result.

Proposition 2. *The performance region $\mathcal{F}$ is a submodular correspondence.*

The correspondence $\mathcal{F}$ is defined for individual resources and can be combined to obtain the performance region for any portfolio of resources in the following fashion. If the portfolio contains K_r units of resource r, then these can be used satisfy any demand that lies in the set $K_r\times\mathcal{F}(r)$. Thus, the performance region of a portfolio of resources, $K=(K_r\in\mathbb{Z}_+: r\in\{0,1\}^n)$, is $\sum_{r\in\{0,1\}^n}K_r\times\mathcal{F}(r)$.

3.3. Capacity portfolio optimization

We next discuss capacity portfolio optimization. This entails attributing costs to resources and costs or quality of service constraints for unmet demand, and then minimizing the overall cost. Our goal is to gain an insight into optimal capacity investment decisions without committing to any specific shortage cost function. To do so, we assign costs to the resources and compare capacity portfolios using their resource costs and performance regions. In particular, we denote the per unit cost of resource $r\in\{0,1\}^n$ as $c(r)$ so that the total resource cost of a portfolio K equals $\sum_{r\in\{0,1\}^n}c(r)K_r$, and then compare any two capacity portfolios K and K' in the following fashion: portfolio K dominates K' if portfolio K has a performance region at least as large as that of K'

and has a resource cost of no greater than that of portfolio K'. If either of these relations holds in a strict manner, then the domination is strong in the sense that portfolio K is strictly preferred to K' in any optimization problem. Using this ranking of portfolios, we obtain the following insight into the optimal portfolio under any objective function.

PROPOSITION 3. *If the resource cost function $c(r)$ is a supermodular function, then there exists an optimal portfolio such that for any pair of resources u, v invested in, if v is more flexible than u, i.e., $v \geq u$, then v can process at most one more product type than u.*

This result proves that for supermodular resource cost functions (which includes the case of affine functions), the ideal resource configuration should consist of resources that are similar to each other in their abilities. For instance, a pair of resources that can process product types $\{1,2\}$ and $\{1,3\}$ is preferred to the pair that processes $\{1\}$ and $\{1,2,3\}$. This result holds for any possible demand realization, and thus does not depend on how one incorporates the demand into the objective function. So, this result remains valid when one considers shortage costs of unmet demand or quality of service constraints, and in this sense is quite powerful. Note that the chaining configurations proposed in Jordan and Graves (1995) comprise resources with similar flexibility, and is consistent with this result.

3.3.1. Capacity portfolio optimization in a three-product example. In this example, we demonstrate the optimal portfolio structure in a parallel newsvendor network with three product types. There are seven types of resources available: three dedicated resources, each capable of processing one type of product; three bi-flexible resources, each capable of processing two product types; and a fully flexible resource that can process all three product types.

As before, each unit of resource $r \in \{0,1\}^3$ costs $c(r)$. While our analysis does not require a formal characterization of the cost of demand shortfall, to relate this setting to a newsvendor setting, we consider the case of linear shortage costs of unmet demand, i.e., there is a per unit demand shortage cost of p_i per unmet demand of product i. The system manager's decision is to select the capacity portfolio before the demand is realized in order to minimize the total system costs. Thus, the optimization problem is

$$\min_K \mathbb{E}_D \left[\sum_{i=1}^{3} p_i (D_i - x_i^*(D)) \right] + \sum_{r \in \{0,1\}^3} c(r) K_r, \tag{3}$$

where $x^*(D)$ is the optimal recourse allocation of the resources to the products and solves

$$\min_{x,y} \sum_{i=1}^{3} p_i (D_i - x_i)$$

$$s.t.\ x_i = \sum_{r\in\{0,1\}^3} y_{i,r} \le D_i \text{ for } i=1,2,3, \tag{4}$$

$$\sum_{i=1}^{3} y_{i,r} \le K_r, \text{ for } r\in\{0,1\}^3 \tag{5}$$

$$y_{i,r} \le K_r r_i, \text{for } i=1,2,3, \text{ and } r\in\{0,1\}^3. \tag{6}$$

In the recourse decision, $y_{i,r}$ refers to the capacity of resource r allocated to product type i and x_i refers to the total capacity allocated to product type i. (4) places the constraint that x_i must be less than the product demand, (5) constrains the capacity allocation from each resource to be less than the number of available units of that resource, and finally (6) ensures that resource r cannot be allocated to product i if it cannot process it, that is, if $r_i = 0$.

We then obtain the following result directly from Proposition 2 for this setting.

Corollary 1. *For a three-product newsvendor setting with a supermodular resource cost function, the optimal solution to* (3) *cannot simultaneously invest in dedicated and fully flexible resources.*

An example of a supermodular resource cost function is one in which the per unit cost of a resource depends only on the number of different products it can process and the marginal cost of increased flexibility (or flexibility premium) is constant or increasing in the number of products. Corollary 1 states that if the resource cost function is supermodular, then an optimal portfolio cannot contain the fully flexible resource and a dedicated resource together. Notice that supermodularity of the resource cost function implies that the cost of one dedicated and one fully flexible resource is no less than that of two appropriately chosen bi-flexible resources, which would in fact have a larger performance region. Specifically, consider a portfolio that consists of the dedicated resource (1,0,0) and the fully flexible resource (1,1,1), and another that consists of the bi-flexible resources (1,1,0) and (1,0,1). Then, we have $c(1,0,0)+c(1,1,1)\ge c(1,1,0)+c(1,0,1)$ using the supermodularity of the resource cost function, and $\mathcal{F}(1,0,0)+\mathcal{F}(1,1,1)\subseteq \mathcal{F}(1,1,0)+\mathcal{F}(1,0,1)$ using the submodularity of the performance region. Thus, the portfolio of bi-flexible resources dominates that containing the dedicated and fully flexible resource, and we obtain that an optimal portfolio may only consist of resources at the same level of flexibility, or at adjacent levels of flexibility.

Notice that this result does not rely on the choice of the shortage penalty function. In particular, we can replace the linear penalty structure for unmet demand with a general function that depends on the demand realized and the consequent total capacity allocated to each product type, and Corollary 1 applies unchanged.

Extending this insight to systems with more than three products, we obtain that if the resource cost function is supermodular, then the optimal portfolio invests in resources that are "close" to each other in the sense that there is no other resource that has a skill-set strictly greater than one resource of the portfolio and strictly less than another. This result is analogous to that in Bassamboo, Randhawa and Van Mieghem (2010). In that paper, the authors focus on newsvendor systems with $n \geq 3$ products with linear penalty costs of unmet demand and constant flexibility premium, and obtain this result by proving that flexibility has diminishing returns.

3.4. Extending the model

We extend the model of Section 3.2 to allow resources to have different levels of productivity. In particular, one unit of a resource can now be used to process different units of various product types. This can be perceived as the resource's efficiency in processing the product types, that is, if one unit of a resource can process more products of one type than that of another resource, then it is said to be more productive in processing that product type.

We now represent a resource by the amount of each product that can be processed per unit, so each unit of a resource $r \in \mathbb{R}^n_+$ processes r_i units of product type i (note that if the resource cannot process product i, then $r_i = 0$). We consider two cases, one in which each unit of a resource can only be used to satisfy the demand of one product type, and the other in which resources can be used to partially satisfy multiple product types simultaneously. Let $\mathcal{G}, \mathcal{H} : \mathbb{R}^n_+ \to \mathcal{P}(\mathbb{R}^n_+)$ denote the correspondence that maps the resources to the demand vectors that they can satisfy for each of these cases, respectively. Then, we can write

$$\mathcal{G}(r) = \{y \in \mathbb{R}^n_+ | \exists p_i \in \{0,1\} \text{ for } 1 \leq i \leq n \text{ such that } \sum_{i=1}^{n} p_i \leq 1 \text{ and } y_i \leq p_i r_i \text{ for } 1 \leq i \leq n\}, \tag{7}$$

$$\mathcal{H}(r) = \{y \in \mathbb{R}^n_+ | \exists 0 \leq p_i \leq 1 \text{ for } 1 \leq i \leq n \text{ such that } \sum_{i=1}^{n} p_i \leq 1 \text{ and } y_i \leq p_i r_i \text{ for } 1 \leq i \leq n\}. \tag{8}$$

Notice that if $r_i \in \{0,1\}$ for all i and y is restricted to the set of integers, then $\mathcal{G}$ reduces to the correspondence $\mathcal{F}$ defined in (2). The correspondence $\mathcal{H}$ generalizes the definition of $\mathcal{G}$ by allowing the resources to be simultaneously used to satisfy demands of different product types.

We then obtain the following result.

Proposition 4. *The performance regions $\mathcal{G}$ and $\mathcal{H}$ are submodular correspondences.*

As in the previous section, we next characterize properties of an optimal capacity portfolio.

Proposition 5. *For a supermodular resource cost function $c(r)$, there exists an optimal capacity portfolio such that for any pair of resources u, v invested in, if $v > u$, i.e., v is either more flexible or more productive or both, then $(v - u)$ has at most one non-zero component.*

This result states that if in an optimal portfolio, a resource v dominates another resource u, and can also process more product types, then it must be the case that resource v can process exactly one more product type than resource u; furthermore, the product types that are common between the two resources must be processed by both at the same rate. This leads to two insights: first, that the optimal portfolio should invest in resources that are close to each other in their skills; and further, that resources that are less flexible must be more productive.

4. Conclusion

This paper generalizes the notion of supermodularity for real-valued functions to correspondences. These submodular correspondences are useful in characterizing optimal flexibility portfolios in newsvendor settings. We prove that the performance region of flexible resources in such settings is a submodular correspondence. When combined with supermodular flexibility cost structures, this yields the result that the optimal flexibility portfolio must consist of resources with similar flexibility. For the general model in which resources may differ in their processing efficiency in addition to their flexibility, we prove that resources that are less flexible must be more efficient in the optimal portfolio.

This paper focuses on static capacity allocation decisions. However, this analysis can also be used in dynamic settings in which the resource allocation decisions in a period have no bearing on future allocations. For instance, if resources are renewed at the end of each period, or if the resources are perishable. In these scenarios, the performance region in each period is identical to that in the static settings, and hence the corresponding results carry through. Note that these results also apply to the resources in newsvendor networks (cf. Van Mieghem and Rudi, 2002).

The static performance region is also applicable when applied to multi-period problems where resources can be dynamically reassigned, as in the case of queueing networks with pre-emption. Here, the demand consists of different types of jobs or customers, each having different processing requirements and the resources do not get consumed during processing. If pre-emption is permitted and causes no losses, i.e., the resource that resumes processing on a job only needs to process the remaining work, then jobs can be transferred from one resource to another seamlessly. In this case, the performance region can be defined as the number of jobs of each type that can be in process at any time instant, and thus our results for the static setting apply. In the case that pre-emption is not practical, allowing for it can still provide a good approximation to the actual system (see the discussion in Atar (2005) for a rigorous justification in queueing systems with large volume), and thus our results may still be applicable.

Unfortunately, when resource allocation decisions made in one period may affect allocations in future periods, as in multi-period inventory models, our analysis does not carry through. This occurs due to the fact that the underlying performance region may now depend on the entire sample path of demand realizations and need no longer be a submodular correspondence. For instance, consider the three product example in Section 3.3.1. In the static setting, the submodularity of the performance region implies that the resource pair of two bi-flexible resources that can process product types 1 and 2, and types 1 and 3, i.e., $(1,1,0)$ and $(1,0,1)$, is preferred to the one resource dedicated to product 1 and one that is fully flexible, i.e., $(1,0,0)$ and $(1,1,1)$. Now consider these two resource configurations but in a two-period setting with the following demand distribution: in the first period, one unit of product 1 is requested, while in the second period, the demand is equally likely to be one unit of product 2 or one unit of product 3. Notice that the resource combination $(1,0,0)$ and $(1,1,1)$ can always fully satisfy the demand by meeting the period one demand using the resource $(1,0,0)$ and meeting the period two demand using the resource $(1,1,1)$, whereas the resource combination $(1,1,0)$ and $(1,0,1)$ may fail to satisfy the demand with 50% chance. Thus, the performance region approach that enumerates all demand scenarios that can be satisfied by the resource is not sufficient to compare these two portfolios of resources, and one must incorporate the actual demand distribution to ascertain which portfolio is better.

References

Atar, R. (2005), ‘Scheduling control for queueing systems with many servers: Asymptotic optimality in heavy traffic’, *Ann. Appl. Probab.* **15**(4), 2606–2650.

Bassamboo, A., Randhawa, R. S. and Van Mieghem, J. A. (2010), ‘Optimal flexibility configurations in newsvendor networks: Going beyond chaining and pairing’, *Management Science* **56**(8), 1285–1303.

Jordan, W. and Graves, S. C. (1995), ‘Principles on the benefits of manufacturing process flexibility’, *Management Science* **41**(4), 577–594.

Simchi-Levi, D., Chen, X. and Bramel, J. (2005), *The logic of logistics: Theory, algorithms, and applications for logistics and supply chain management*, Springer-Verlag, New York, NY.

Van Mieghem, J. A. and Rudi, N. (2002), ‘Newsvendor networks: Inventory management and capacity investment with discretionary activities’, *Manufacturing and Service Operations Management* **4**(4), 313–335.

Appendix. Proofs

We provide proofs for Propositions 1, 4, and 5. Propositions 2 and 3 are special cases of Proposition 4 and 5, respectively, and we omit their proofs.

Proof of Proposition 1. Part 1 follows immediately from the definition of supermodularity.

To prove part 2, it suffices to show that $Co(A)+Co(B)=Co(A+B)$ for any $A,B\in\mathcal{P}(\mathbb{R}^n)$, where $Co(X)$ denotes the convex hull of $X=A,B$. It is easy to see that $Co(A+B)\subseteq Co(A)+Co(B)$. To complete the proof, we show that $Co(A)+Co(B)\subseteq Co(A+B)$. Let $x\in Co(A)+Co(B)$ such that $x=a+b$, $a\in Co(A)$, and $b\in Co(B)$. That is, there exists m and $a_1,a_2,\cdots,a_m\in A$ and $\alpha\in\mathbb{R}^m_+$ with $\sum_{i=1}^m\alpha_i=1$ such that $\sum_{i=1}^m\alpha_i a_i=a$ and n and $b_1,b_2,\cdots,b_n\in B$ and $\beta\in\mathbb{R}^n_+$ with $\sum_{i=1}^n\beta_i=1$ such that $\sum_{i=1}^n\beta_i b_i=b$. We also have $a_i+b_j\in A+B$ for all $1\le i\le m$ and $1\le j\le n$. It then follows that $x=\sum_{i,j}\alpha_i\beta_j(a_i+b_j)\in Co(A+B)$. ■

Proof of Proposition 4. We prove that $\mathcal{G}$ is a submodular correspondence. Noting that $\mathcal{H}$ is the convex hull of $\mathcal{G}$, applying Proposition 1 it then follows that $\mathcal{H}$ is a submodular correspondence. Denote e^i as the unit vector with its i^{th} component equal to 1, and all other components equal to zero 0. Then, noting that a resource can only be used to satisfy the demand for one product, we can write $\mathcal{G}(x)=\cup_{i=1}^n\{z|z\le x_ie^i\}$ for $x\in\mathbb{R}^n_+$ and similarly, for $y\in\mathbb{R}^n_+$, we have $\mathcal{G}(y)=\cup_{i=1}^n\{z|z\le y_ie^i\}$. This implies that

$$\begin{aligned}\mathcal{G}(x)+\mathcal{G}(y)&=\left(\cup_{i=1}^n\{z|z\le(x_i+y_i)e^i\}\right)\cup\left(\cup_{i=1}^n\cup_{j=1,j\neq i}^n\{z|z\le x_ie^i+y_je^j\}\right)\\ \mathcal{G}(x\vee y)+\mathcal{G}(x\wedge y)&=\left(\cup_{i=1}^n\{z|z\le(x_i+y_i)e^i\}\right)\cup\left(\cup_{i=1}^n\cup_{j=1,j\neq i}^n\{z|z\le(x_i\vee y_i)e^i+(x_j\wedge y_j)e^j\}\right).\end{aligned}$$

To prove that $\mathcal{G}(x\vee y)+\mathcal{G}(x\wedge y)\subseteq\mathcal{G}(x)+\mathcal{G}(y)$, it suffices to prove $\{z|z\le(x_i\vee y_i)e^i+(x_j\wedge y_j)e^j\}\subseteq\{z|z\le x_ie^i+y_je^j\}\cup\{z|z\le y_ie^i+x_je^j\}$. Without loss of generality, we can assume that $x_i\vee y_i=x_i$. Then, $\{z|z\le(x_i\vee y_i)e^i+(x_j\wedge y_j)e^j\}\subseteq\{z|z\le x_ie^i+y_je^j\}$ because $x_j\wedge y_j\le y_j$, and the result follows. ■

Proof of Proposition 5. Let K^* be the optimal portfolio with the smallest $\sum_{r\in\mathbb{R}^n_+}(e'r)^2K_r$, where $e=(1,1,\cdots,1)$. We prove that K^* has the stated property arguing via contradiction. Suppose K^* invests in u and v such that $(v-u)$ has at least two non-zero components. Then, pick any j such that $v_j-u_j>0$ and define $\delta=(v_j-u_j)e^j$, where e^j is the unit vector with its j^{th} component equal to 1, and all other components equal to zero 0. Consider now a new portfolio K' constructed from K^* by replacing one unit of resource u and one unit of resource v with one unit of resource $s\equiv u+\delta$ and one unit of resource $t\equiv v-\delta=u+v-s$. Notice that $s\vee t=v$ and $s\wedge t=u$. Proposition 4 implies that K' has a performance region that is at least as large as K^*, and further because $c(r)$ is a supermodular function, the cost of the capacity portfolio K is weakly lower than that of K^*. If the cost of the capacity portfolio K is strictly lower than that of K^*, we obtain a contradiction to the optimality of K^*. Otherwise, both the capacity portfolios K^* and K have the same cost, and then noting that $e'u<e's<e'v$ and $e'u<e't<e'v$, and $e'u+e'v=e's+e't$, we find that $(e'u)^2+(e'v)^2>(e's)^2+(e't)^2$. This implies that $\sum_{r\in\mathbb{R}^n_+}(e'r)^2K'_r<\sum_{r\in\mathbb{R}^n_+}(e'r)^2K^*_r$, and we obtain a contradiction to the fact that K^* minimizes $\sum_{r\in\mathbb{R}^n_+}(e'r)^2K_r$. ■