Fast and frugal heuristics for portfolio decisions with positive project interactions

Ian N. Durbach*[1,2,3], Simón Algorta[4], Dieudonné Kabongo Kantu[1,3,5], Konstantinos V. Katsikopoulos[6], and Özgür Şimşek[7]

[1]Centre for Statistics in Ecology, the Environment and Conservation, Department of Statistical Sciences, University of Cape Town, South Africa

[2]Centre for Research into Ecological and Environmental Modelling, University of St Andrews, UK

[3]African Institute for Mathematical Sciences, South Africa

[4]Faculty of Mechanical Engineering and Transport Systems, Technische Universität Berlin, Germany

[5]Ipsos Laboratories, Cape Town, South Africa

[6]Centre for Operational Research, Management Science and Information Systems, University of Southampton Business School, UK

[7]Department of Computer Science, University of Bath, UK

Abstract

We consider portfolio decision problems with positive interactions between projects. Exact solutions to this problem require that all interactions are assessed, requiring time, expertise and effort that may not always be available. We develop and test a number of fast and frugal heuristics – psychologically plausible models that limit the number of assessments to be made and combine these in computationally simple ways – for portfolio decisions. The proposed "add-the-best" family of heuristics constructs a portfolio by iteratively adding a project that is best in a greedy sense, with various definitions of "best". We present analytical results showing that information savings achievable by heuristics can be considerable; a simulation experiment showing that portfolios selected by heuristics can be close to optimal under certain conditions; and a behavioral laboratory experiment demonstrating that choices are often consistent with the use of heuristics. Add-the-best heuristics combine descriptive plausibility with effort-accuracy trade-offs that make them potentially attractive for prescriptive use.

Keywords: Decision making; decision analysis; portfolio selection; heuristics; behavioural decision making

*Corresponding author: ian.durbach@uct.ac.za

1 Introduction

Portfolio decisions involve selecting a subset of alternatives or "projects" that together maximize some measure of value, subject to resource constraints (Salo et al., 2011). Examples include capital investment (Kleinmuntz, 2007; Airoldi and Morton, 2011), R&D project selection (Phillips and Bana e Costa, 2007; Jung and Seo, 2010; Arratia et al., 2016; Liesiö and Salo, 2012; Jang, 2019), maintenance planning (Mild et al., 2015), and windfarm location (Cranmer et al., 2018). This paper considers portfolio problems in which benefits and costs are not necessarily additive: some projects may interact with one another.

Exact solutions to this problem require that all project interactions are assessed, and the time and effort involved in this can be considerable. As the starting point for this paper we take the view that in some problems project interactions can only be assessed by consulting a human decision maker or expert, and that sometimes the number of interactions will be too large for the assessment of all of them to be feasible. The purpose of this paper is to propose several heuristics that limit the number of assessments that are made and thus may be suitable for portfolio decision problems in which the complete assessment of interactions is not an option. We evaluate these heuristics in terms of how many assessments they save, and how close their portfolio values are to the theoretical optimal value that would be achieved if all interactions were known and exact methods used. We also use a behavioral laboratory experiment to provide evidence of behaviour that is consistent with using some of the proposed heuristics.

We draw a distinction between our heuristics and those developed in the optimization literature, where the problem above has been extensively studied for decades, either in its interaction-free version as the standard knapsack problem or, with some restrictions (value interactions involving pairs of projects only) as the quadratic knapsack problem. Exact algorithms (pseudo-polynomial in the standard case), efficient approximations, and numerous computational heuristics have been developed for both problems (Pisinger, 2007). These require all interactions to be assessed upfront and their goal is to limit the amount of computation time required to solve the problem. This is important when the number of projects is very large, but less relevant when projects number in the tens or hundreds, as is typically the case for portfolio problems in which decision support is provided (see e.g. applications reported in Salo et al. (2011)). In these cases using a computational heuristic is inappropriate – if all interactions can be assessed then an exact method should be used. The heuristics we propose address a different kind of time- and effortsaving to computational heuristics – time and effort in assessment – and are in the tradition of so-called fast and frugal heuristics (Gigerenzer et al., 1999) or psychological heuristics (Keller and Katsikopoulos, 2016), which use limited information and process this information in computationally simple ways e.g. elimination-by-aspects Tversky (1972), take-the-best (Gigerenzer and Goldstein, 1996). These heuristics are typically not normative, but invoke bounded rationality arguments to argue for both potential prescriptive use (if environments in which cases good performance is obtained are known) and descriptive plausibility (Gigerenzer and Goldstein, 1996). Different heuristics may of course vary in the degree to which they emphasise prescriptive or descriptive aspects (Todd, 2007; Katsikopoulos et al., 2018).

Our heuristics construct a portfolio by iteratively adding a project that is best in a greedy (i.e. locally optimal) sense. Sharing this common structure, we collectively call them the *add-the-best* family of heuristics. For example, in a computationally demanding version of add-the-best, the "best" project is the one whose selection leads to the largest immediate increase in portfolio value, including the value added by project interactions. In computationally simpler heuristics, a best project is again one which leads to the largest immediate increase in portfolio value, but this is now calculated without considering interactions. Add-the-best heuristics are conceptually closely related to single-cue heuristics that make decisions using a single piece of information; in cases where this single piece of information does not discriminate among the projects, the heuristic decides randomly (Hogarth and Karelaia, 2005).

The primary goal of our paper is to extend fast and frugal heuristics, which have been extensively studied in traditional choice problems, to portfolio decision making involving project interactions. We find that, in contrast to choice problems, where simple heuristics often perform unexpectedly well (e.g. Hogarth and Karelaia, 2005; Todd, 2007), it is much harder to strike a balance between frugal information use and good performance in portfolio problems. Our main contribution is to develop two heuristics called *Added Value* and *Unit Value with Synergy* that achieve this balance, returning portfolios that are competitive with those obtained by exact methods while limiting the number of assessments to potentially manageable levels. These heuristics combine descriptive plausibility with effort-accuracy trade-offs that make them potentially attractive for prescriptive use in cases where complete assessement of interactions is not feasible.

2 Portfolio decision making

Stummer and Heidenberger (2003) describe the formulation of the portfolio decision problem with interactions, whose goal is to decide which projects to select from a set of candidates $\{P_1, \ldots, P_J\}$, so as to maximize the overall value of the portfolio subject to budget and any other constraints. Interactions between projects are modelled by defining interaction subsets $\mathcal{A}_k$ containing those projects making up interaction $k = 1, \ldots, K$. A set $\mathcal{A}_k$ is defined for each subset of projects whose total value or cost is not simply the sum of their individual values and costs. Overall portfolio value is given by

$$V(\mathbf{z}) = V(z_1, \ldots, z_J) = \sum_{j=1}^{J} b_j z_j + \sum_{k=1}^{K} B_k g_k \tag{1}$$

where b_j is the individual value of project P_j if implemented on its own, $z_j = 1$ if project P_j is selected ($z_j = 0$ otherwise), B_k is the incremental change in value if all of the projects in interaction subset $\mathcal{A}_k$ are included in the portfolio, and $g_k = 1$ if all projects in interaction subset $\mathcal{A}_k$ are selected ($g_k = 0$ otherwise). This is to be maximized, subject to the budget constraint

$$C(\mathbf{z}) = C(z_1, \ldots, z_J) = \sum_{j=1}^{J} c_j z_j + \sum_{k=1}^{K} C_k g_k \leq \zeta \tag{2}$$

where c_j is the individual cost of project P_j if implemented on its own, C_k is the incremental change in cost if all of the projects in interaction subset $\mathcal{A}_k$ are included, ζ is the total budget, and z_j and g_k are as defined previously. We restrict ourselves to cases where interactions are expressed as positive increases in value ($B_k \geq 0$, $C_k = 0, \forall k$). For convenience, we sometimes refer to the budget in relative terms, as a proportion of the sum of individual costs i.e. $\zeta / \sum_{j=1}^{J} c_j$.

The problem above can be formulated as an integer linear program using auxiliary constraints to define the g_k, and solved using standard techniques (Stummer and Heidenberger, 2003), provided that all interactions are known. Many extensions have been proposed to treat different kinds of interactions (Liesiö et al., 2007; Liesiö, 2014; Barbati et al., 2018; Cranmer et al., 2018; Vilkkumaa et al., 2018; Korotkov and Wu, 2020). These too require the complete enumeration of interactions in order to compute the optimal portfolio and so are not discussed further here. Methods are available for cases where the coefficients in (1) or (2) e.g. those capturing interaction values and costs, are imprecisely known. These either integrate out uncertainty to maximize some combination of expected value and risk (e.g. Hassanzadeh et al., 2014; Jang, 2019), or identify sets of potentially optimal portfolios and provide robustness diagnostics on these, rather than select a single portfolio (e.g. Lourenco et al., 2012; Baker et al., 2020). All methods still require the assessment of all interactions, even though these can be imprecise.

Heuristics (Tversky, 1972; Gigerenzer and Goldstein, 1996; Katsikopoulos, 2011) have been extensively studied for traditional (one-out-of-n) choice problems. Findings indicate with reasonable confidence that (a) psychologically plausible heuristics can offer outcomes that are competitive with theoretically optimal models under reasonably well-known conditions (Hogarth and Karelaia, 2005; Todd, 2007; Baucells et al., 2008; Buckmann and Şimşek, 2017; Katsikopoulos et al., 2018), (b) some of these conditions often occur in real-world contexts (Şimşek, 2013), and (c) decision makers use heuristics, particularly when time pressure or the cost of gathering information is high (Ford et al., 1989; Bröder and Newell, 2008).

Very little equivalent work exists for portfolio problems (Fasolo et al., 2011; Schiffels et al., 2018), particularly for (a) and (b) above and even more so when project interactions are involved. Keisler (2004, 2008) implemented a portfolio heuristic that adds projects in order of their value-to-cost ratios (our *Unit Value* heuristic). The focus of the paper was on the value of gathering additional information about project values and costs when these were initially uncertain, so that heuristic performance (relative to an optimal solution) was not assessed. Interactions were also not included. A later working paper (Keisler, 2005) included interactions, but again focused on improvements in portfolio value achieved by gathering additional information (this time about the interactions themselves). All possible portfolios were enumerated, so no selection heuristics were used.

The few behavioral studies to date have suggested that many decision makers use some form of heuristic reasoning when solving portfolio problems. When solving standard knapsack problems without interactions, untrained participants commonly selected projects by sorting on their value-to-cost ratios or, to a lesser extent, on their costs or value-to-cost differences (Schiffels et al., 2018; Pape et al., 2019), with evidence of multiple heuristic use over the course of the experiment (Schiffels et al., 2018) and a bias towards selecting low-cost projects (Pape et al., 2019). Phillips and Bana e Costa (2007) report that 23 out of 28 companies used judgments such as ranking projects by expected benefit and adding these until reaching a budget limit (our *Highest Value* heuristic) to prioritize drug development, a higher proportion than achieved by any mathematical model. Langholtz and colleagues show both novice and experts use heuristics that they group into "solve-and-schedule" and "consume-and-check" strategies to allocate resources across projects (Langholtz et al., 1993, 1997; Ball et al., 1998; Langholtz et al., 2002). Solve-and-schedule strategies start by setting a total objective function value and then allocate resources across projects so that this value is achieved. Consume-and-check strategies make a sequence of related decisions about which resource to consume "next", at each stage checking on remaining resources and constraint violations. In a key experiment participants decided how to allocate their time and money to consume a maximum number of meals of either restaurant or home-cooked "types". A solve-and-schedule approach decides on the total number of meals and then searches for ways to allocate these between meal types without violating constraints, while consume-and-check asks only whether the next meal should be from a restaurant or home-cooked.

These descriptive studies motivate and inform our work but tend to employ decision problems that support their aim of inferring descriptive detail, an aim quite different to our own. For example, Langholtz et al. (1997) use resource allocation problems where there are only two types of projects, people can consume many of each, and each project type shares the same benefit and cost values. This simplifies the context and makes solving to optimality possible (using graphical methods) even if it is unlikely. The problem we address involves selecting a best subset from a discrete set of projects, all of which differ in terms of benefits and costs. Each project can be selected once or not at all. Solve-and-schedule strategies are unlikely in contexts like these, because the "solve" step requires assessing a desired overall portfolio value from dozens of projects with different costs, benefits, and interactions. Adding projects sequentially, which is by definition a "consume-and-check" heuristic, would seem to be the rule (see also Rieskamp et al. (2003)). There is no simple mapping of consume-and-check heuristics to the heuristics we propose. Fasolo et al. (2011) point out that the resource allocation and best-subset selection formulations are only the same "where projects are associated with particular organisational subunits (i.e. projects can be partitioned into subsets of projects which 'belong' to particular subunits)", which is not the case here. Finally, interactions are not considered, and all project information is known beforehand. In contrast our focus is on interactions, which individuals must assess as they go.

3 Proposed fast and frugal portfolio heuristics

In this section we propose a family of fast and frugal heuristics for selecting portfolios. A numerical example illustrating each heuristic is given in Appendix A. The heuristics are frugal in that they do not use all of the available information, and fast because they integrate the information in simple ways to decide which project to include next, and when to stop. All except one uses a single well-defined criterion in adding projects to the portfolio, extending single-cue heuristics developed for simpler decision problems (such as choice and comparison) into the domain of portfolio selection problems.

Our heuristics construct portfolios by sequentially adding projects, excluding those additions that would, if implemented, violate budget or other logical (e.g. project interaction) constraints[1]. We specify a stopping rule by which portfolio construction terminates after a user-specified number of consecutive constraint violations. Note that setting this number suitably large guarantees an exhaustive search through the list of projects. We call the proposed family of heuristics *Add-the-best*.

Add-the-best A family of heuristics for portfolio selection. Starting with an empty set of selected projects, at each stage the heuristics evaluate those projects not yet added to the portfolio. Evaluation is independent and over a single well-defined criterion. The project that has the highest value on this criterion is added to the portfolio provided its addition does not violate budget constraints. Ties are broken randomly. Individual heuristics in the family differ on the criterion they use in evaluating candidate projects. The process terminates after a user-specified consecutive violations of the budget constraint or when no projects remain to be considered.

We first define three heuristics that do not use project interactions at all. While these heuristics may appear excessively simple, there is evidence that they are used in real-world portfolio decision making (Phillips and Bana e Costa, 2007; Schiffels et al., 2018) and they provide a useful starting point for our study by allowing us to measure the impact of ignoring interaction information on overall portfolio value.

Highest Value Adds projects in descending order of their values.

[1] Constraints on project combinations are most easily handled in this way i.e. as a veto, but it is also possible to modify add-the-best heuristics so that, for example, if an already-included project is repeatedly involved in interaction violations that prevent the addition of otherwise good projects, then that project is removed.

Lowest Cost Adds projects in ascending order of their costs.

Unit Value Adds projects in descending order of their value-to-cost ratios. Values are based on individual project values only.

To these three heuristics we add a fourth that makes use of dominance relationships. In this case, the criterion for "best" is simply that the project is not dominated by any project that remains outside the portfolio (in the sense of having both a lower value and higher cost e.g. Lourenco et al. (2012)) .

Pareto This heuristic adds a randomly chosen project provided it is within budget and does not have both a lower value and higher cost that any project not already in the portfolio.

We base dominance assessments on individual values and costs only, although other information could also be used. For example, dominance across multiple attributes is easily assessed and thus the heuristic extends easily to a multi-attribute context. Importantly, we consider dominance relations only between projects that are not already part of the portfolio. Our motivation is that while we do not want to add a project that is unambiguously worse than another candidate project, portfolios may well be improved by the addition of projects that are dominated by one of the already selected projects. For example, in cases where a single project dominates all others we would still want to add further projects until the budget is reached. The *Pareto* heuristic can pick many different sets of projects because it involves, at each step, a random selection from the set of non-dominated candidates.

The four heuristics above ignore all information about project interactions. Our next heuristic uses binary information indicating whether a project is involved in any positive interaction, without evaluating the number or magnitude of these interactions, and uses this information to preferentially select projects that are involved in positive interactions. This provides a bridge to heuristics that make use of the magnitude of project interactions.

Unit value with Synergy Identifies all projects that are involved in at least one positive interaction. Adds projects from this set using the *Unit Value* heuristic i.e. in descending order of their value-to-cost ratios, with values based on individual project values only. Once this set has been exhausted, adds projects from outside the set, again using *Unit Value*.

Our remaining heuristics make use of quantitative information about interactions between projects. These remain greedy (projects are added to the portfolio one at a time) and naive (eligible projects are evaluated independently), and differ from one another depending on whether they consider *all* interaction subsets or restrict themselves to a subset of the interactions. We first consider a heuristic that uses all interactions:

Added Value This heuristic adds the project whose selection would lead to the largest increase in overall portfolio value per unit cost. The incremental benefit includes the individual value of the project, as well as the value of all interaction subsets that would be completed if the project were to be added.

At each step, *Added Value* must search over all interaction subsets that are not already active, each time assessing whether adding a particular project would complete any of the interaction subsets. More frugal heuristics do not search all interaction sets, but only those that fulfill some additional criteria. We list three such heuristics below – although only the first has an intuitive appeal, the others allow us to examine the sensitivity of heuristics to how the shortlist of interaction subsets is constructed.

Added Value Most This heuristic only considers interaction subsets that involve the project that currently contributes the most to portfolio value. When assessing which project contributes most, the contribution of each project already in the portfolio is defined as the decrease in portfolio value that would be experienced if the project was removed. This includes the marginal value of the project as well as the value of any complete interaction subsets the project belongs to. The incremental benefit of a project not already in the portfolio is the sum of its individual value and the value of any interaction subsets involving the most valuable project that would be completed by the addition of the project to the portfolio.

Added Value Least This heuristic is defined as *Added Value Most* except that it considers only interaction subsets that involve the project that currently contributes the *least* to portfolio value.

Added Value Random This heuristic randomly chooses one of the projects already in the portfolio and considers only the interaction subsets that involve this project.

4 Analytical results on information requirements

Exact methods require the assessment of all m-way interactions up to order M. Assuming that M is somehow known, this equates to $\sum_{m=2}^{M}\binom{J}{m}$ interactions. While many of these interactions could easily be ruled out by statements such as "project X does not interact with any other project", the number of interactions provides a useful baseline for comparison with heuristics.

How much information do the add-the-best heuristics use? Let $P_{(s)}$ denote the s-th project added, and $\mathcal{J}_s^*$ denote the set of $J-s$ projects remaining in contention after s projects have been included. We call projects that have not yet been included in the portfolio 'candidate' projects, and those that have been included 'existing' projects.

The number of m-way interactions assessed by *Added Value* can be calculated as follows. No m-way interactions need be assessed until $m-1$ projects are already in the portfolio. At step $s \in \{m-1,\ldots,J-1\}$ there are s projects in the portfolio and $J-s$ candidates. The only *new* m-way interactions that need to be assessed involve (a) the most recently added project $P_{(s)}$, (b) a candidate project $P_j \in \mathcal{J}_s^*$, and (c) $m-2$ other existing projects drawn from $\{P_{(1)},\ldots,P_{(s-1)}\}$. All m-way interactions that do not involve the most recently added project will have already been assessed in previous iterations. There are $J-s$ candidate projects and $\binom{s-1}{m-2}$ ways of arranging the other existing projects in part (c); the number of assessments that *Added Value* needs to do is given by the product $\binom{s-1}{m-2}(J-s)$.

The *Added Value Most* heuristic assesses only a subset of these interactions; those that involve, at a particular step s, the project that contributes most to the portfolio at that time, called the "most valued project" or MVP. The number of new interactions to assess thus depends on whether or not the MVP has changed. Bounds are easily calculated – the upper bound, obtained when the MVP changes at every step, is the number of assessments *Added Value* needs; while the lower bound is obtained as $\binom{s-2}{m-3}(J-s)$, for $m \geq 3$ if the MVP never changes. The same bounds apply to *Added Value Least* and *Added Value Random* heuristics.

The *Added Value* heuristic requires only a small fraction of the assessments required by a full optimization approach, provided that the constructed portfolio contains relatively few projects as a proportion of the total available (Figure 1). As the number of projects that can be selected is almost entirely a function of the available budget, this means that heuristics are relatively more frugal when budgets are limited. If the final portfolio contains 10 out of the 50 available projects, *Added Value* requires 445 (36%) of 1225 two-way, 1920 (10%) of 19600 three-way, 5010

(2%) of 230300 four-way, and 8652 (0.4%) of 2118760 five-way interactions. The more restrictive *Added Value Most* requires a minimum of 49 (4%) of 1225 two-way, 396 (2%) of 19600 three-way, 1524 (0.7%) of 230300 four-way, and 3486 (0.2%) of 2118760 five-way interactions.

The relative reduction from what is required by an optimal model is substantial, particularly with small budgets, but in absolute terms the number of assessments needed by *Added Value* remains large. Practical applications of the heuristic may depend on finding alternate ways of directly estimating the marginal increase in portfolio value, or else ignoring higher-order interactions.

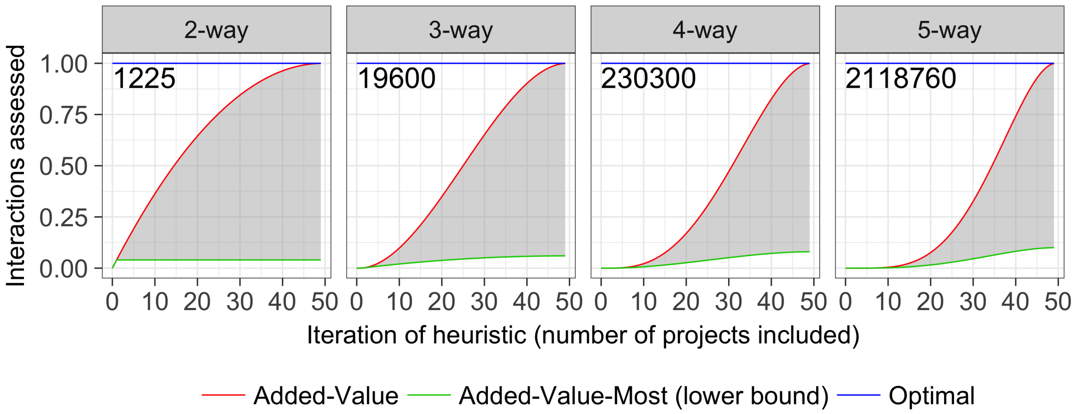

Figure 1: Cumulative number of m-way interactions that need to be assessed by the add-the-best heuristics, expressed as a proportion of the total number of possible interactions for $J = 50$ projects and $m \in \{2, 3, 4, 5\}$. The grey shaded area indicates the lower and upper bounds of the *Added Value Most* heuristic. The total number of interactions i.e. $\binom{50}{m}$) is indicated in the top left corner of each panel). Note that full optimization of portfolio value requires all interactions to be assessed.

The number of assessments required by the *Unit Value with Synergy* heuristic is difficult to specify analytically because it depends on the assessment process used. The heuristic requires only that projects that do not interact at all are removed from consideration. At best this requires at most J questions of the form “does this project have any interactions with any project (or combinations of projects)?” These assessments are of a kind that are not directly comparable with the assessments used by other heuristics. It is also unclear if and under what conditions decision makers can reliably answer these questions, an issue we revisit in Section 7. At worst the heuristic requires the decision maker to assess whether each of the $\sum_{m=2}^{M} \binom{J}{m}$ possible interactions exist, which is certainly impossible. In reality this worst case is highly unlikely because establishing one interaction immediately makes many others redundant, but it is sufficient to demonstrate the challenges in establishing information requirements. Following the removal of non-interacting projects the *Unit Value with Synergy* heuristic applies the *Unit Value* heuristic, which even over the full set of projects is extremely frugal, as are the other heuristics that ignore interactions, *Highest Value*, and *Lowest Cost*. However, as we show in the next section, applying any heuristics ignoring project interactions in an unknown context would seem to require accepting a very high probability of selecting a poor portfolio.

5 Simulation-based comparison of heuristic and optimal portfolios

In previous sections we proposed a number of fast and frugal heuristics for portfolio selection, and showed that these have relatively low information requirements. In this section we evaluate the ability of these heuristics to achieve overall portfolio values comparable with those obtained by optimal portfolios. Our simulation structure consists of (a) generating a number of projects and their individual values and costs, (b) creating interdependencies between the projects, (c) defining the incremental values and costs associated with each of the interaction subsets, (d) running optimal and fast and frugal portfolio selection models, and (e) comparing the values obtained from fast and frugal and optimal portfolios. Simulations were written and analyzed in R 3.6.0 using packages `Rglpk` (Theussl and Hornik, 2019) and `ggplot2` (Wickham, 2016). All code and results are available at `https://github.com/iandurbach/portfolio-heuristics`.

5.1 Simulation study design

5.1.1 Generating individual values and costs

The problem context is defined by the number of projects J, the individual values b_j and costs c_j associated with each project P_j, and the total budget ζ. We simulated problems involving $J = 50$ projects. Individual project values were generated to be either uniform ($b_j \sim U[0.5, 5]$), positively skewed ($b_j \sim Gamma(0.5, 2)$), or negatively skewed ($b_j^* \sim Gamma(0.5, 2)$; $b_j = \max_j b_j^* - b_j + 0.1$). Project costs were generated as $c_j = a_j b_j$, where $a_j \sim U[80, 120]$; the scaling of a_j relative to b_j is unimportant, since we use only one benefit and cost attribute. Generating values and costs in this way means that value *per unit cost* are, on average, uncorrelated with value and weakly negatively correlated with cost (uniform: -0.2; skewed: -0.1). We varied the available budget ζ by choosing the proportion $\zeta / \sum_{j=1}^{J} c_j$ to lie between 0.1 and 0.9 in increments of 0.1. Note that if $\zeta / \sum_{j=1}^{J} c_j = 1$ then all projects can be selected.

5.1.2 Creating interactions between projects

In the following we describe two ways of constructing subsets of interacting projects, which we term *random* and *nested* respectively. Both start by selecting $J^+ \leq J$ projects to create a set of projects $\mathcal{J}^+$ from which interdependencies will be drawn. Projects are selected either with selection probabilities (a) equal across projects, (b) directly proportional to their value-to-cost ratio b_j/c_j, in which case projects that are individually better are more likely to be involved in positive interactions, (c) inversely proportional to b_j/c_j, in which case worse projects are more likely to be involved in interactions. This is a simulation parameter, with conditions (b) and (c) expected to help and hinder heuristics respectively.

Random interactions have no structure linking lower- and higher-order interaction subsets. Each interaction subset is obtained by randomly sampling the required number of projects from $\mathcal{J}^+$, independent of any other interaction subset. With nested interactions, a low-order interaction subset (one containing relatively few projects) is generated by sampling the required number of projects *from one of the already-generated higher-order interaction subsets*, rather than from $\mathcal{J}^+$. For example, in our study we set $J^+ = 10$ and generated two interaction subsets involving five projects, six subsets of four projects, eight subsets of three projects, and ten subsets of two projects. We begin by generating the two highest-order subsets by randomly selecting five projects from the ten in $\mathcal{J}^+$, twice. To generate each of the fourth-order interactions, we randomly select one of the fifth-order interaction subsets and randomly select four projects from this subset. To generate each third-order interaction we randomly select one of the fourth-order interaction subsets and randomly select three projects from this subset. We continue in this fashion until all interactions have been generated.

5.1.3 Computing values and costs of interactions

Our study employs only positive interactions expressed through increases in benefits if certain combinations of projects are selected. We set the incremental benefit of completing interaction subset $\mathcal{A}_k^+$ to be a proportion γ of the sum of the values of projects in $\mathcal{A}_k^+$ i.e. $B_k = \gamma \sum_{j \in \mathcal{A}_k^+} b_j$, with $\gamma \in \{0, 0.5, 1\}$ a parameter of the simulation. Higher-value projects thus result in interactions with higher absolute values, although as these projects also tend to cost more lower-value projects may still be preferred per unit cost. We chose values of γ so that interactions contribute a substantial proportion of the overall value of the optimal portfolio, on a trial-and-error basis. With $\gamma = 0.5$, interactions contribute on average between 22% (at high budgets, $\zeta = 0.9 \sum_{j=1}^{J} c_j$)

and 48% ($\zeta = 0.1 \sum_{j=1}^{J} c_j$) of overall portfolio value. With $\gamma = 1$ these percentages rise to 36% and 65% respectively. Our motivation here is to avoid making overly favourable claims for those heuristics that ignore interactions between projects.

5.1.4 Running portfolio selection models

The optimal portfolio is found by maximizing (1) subject to the budget constraint (2), using the approach in Stummer and Heidenberger (2003). We implemented all nine heuristics described in Section 3, stopping after receiving three budget violations. We also computed (a) the mean value over 100 random feasible portfolios, constructed by randomly adding one of the remaining projects subject to budget constraints, and (b) the value of the worst-case or 'nadir" portfolio, obtained by *minimizing* the objective function in Section 1 subject to the same constraints plus an additional one that forces projects to be chosen until at least 95% of the budget ζ has been spent. Random portfolio construction can be considered fast and frugal, as it terminates in a small number of steps and requires little information, but it is also 'dumb', in the sense that it exploits no information about the projects themselves. It therefore seems a reasonable basis for judging the performance of any other heuristic. Values of the nadir portfolio are shown largely so that the reader can compare these with what is achieved with a random selection.

5.1.5 Comparing results

From each simulation run we obtain the value of the portfolio selected by each of the heuristics, as well as the value of the optimal portfolio. We show performance both in absolute terms, i.e. the values of the portfolios, and in a standardized form in which portfolio values are normalized relative to the optimal portfolio, which is assigned a value of 100.

5.2 Results

The *Added Value* and *Unit Value with Synergy* heuristics perform well across a range of simulated contexts, and offer close to optimal performance with moderate-or-larger budgets (Figure 2). Once the budget is 30% of total cost, the *Added Value* and *Unit Value with Synergy* heuristics achieves 85% and 80% of the available gains respectively. The good performance of the *Unit Value with Synergy* heuristic suggests that quantitative information is not strictly necessary for good performance – knowing only about the *presence* of interactions can improve performance substantially.

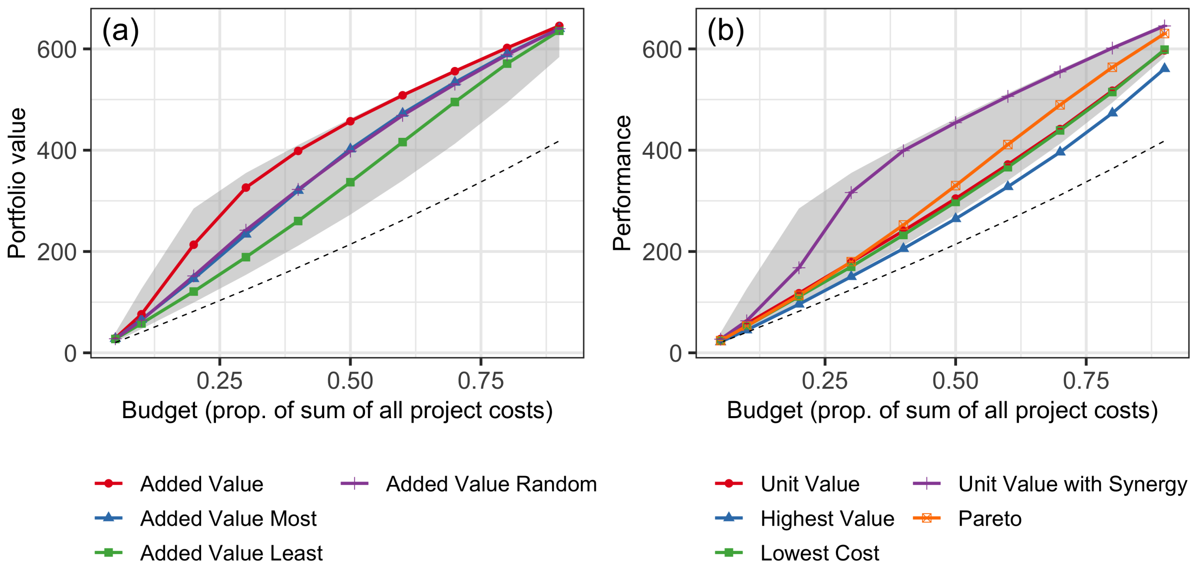

Figure 2: Mean values of portfolios selected by fast and frugal portfolio heuristics under different budget constraints. Panel (a) shows heuristics that consider quantitative project interactions; panel (b) shows heuristics that do not. Confidence intervals around these means are neglible (smaller than the symbols used to plot the means). The grey polygon plots the envelope between the value of the optimal portfolio and the mean value returned by a random selection of projects, which we consider a useful lower bound for benchmarking performance. The dashed line denotes the value of the nadir portfolio.

It is important that all interactions are assessed, as both *Added Value* and *Unit Value with Synergy* do. If not, performance worsens considerably. The set of heuristics *Added Value Most*, *Added Value Least* and *Added Value Random* offer large improvements over randomly selected portfolios but perform substantially worse than *Added Value* or *Unit Value with Synergy*. There are no material differences between the *Added Value Random* heuristic and the *Added Value Most* heuristic over the entire budget range, while as the budget increases the *Added Value Least* heuristic performs substantially worse than the other two. Of the second set of heuristics shown in Figure 2b, those that do not consider interactions between projects at all perform on the whole substantially worse, and cannot in general be recommended as selection strategies. The *Highest Value* heuristic performs worse than *Unit Value* and *Lowest Cost* because project values are highly correlated with project costs, so fewer projects are added before the budget is exceeded and interactions are less likely. The poor performance of *Unit Value* is determined by the magnitude of our simulated interactions, but remains poor even in the smaller of our conditions (Figure 3).

The performance of *Added Value* and *Unit Value with Synergy* at very low budget levels (10% of total cost) is worse when interactions are nested than when they are random (Figure 4).

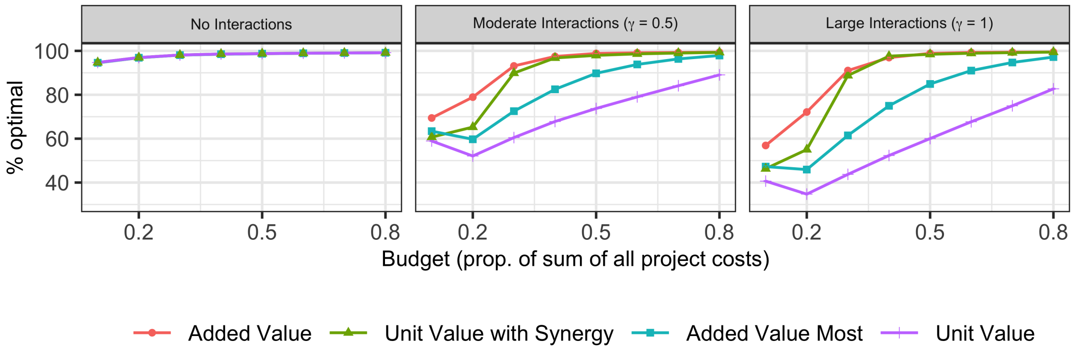

Figure 3: Relative performance of add-the-best variants for different project interaction magnitudes. Projects making up an interaction subset each have individual project values, and hence a sum exists for the interaction subset. The γ parameter indicates the proportion of this sum that is awarded when the entire interaction subset is selected.

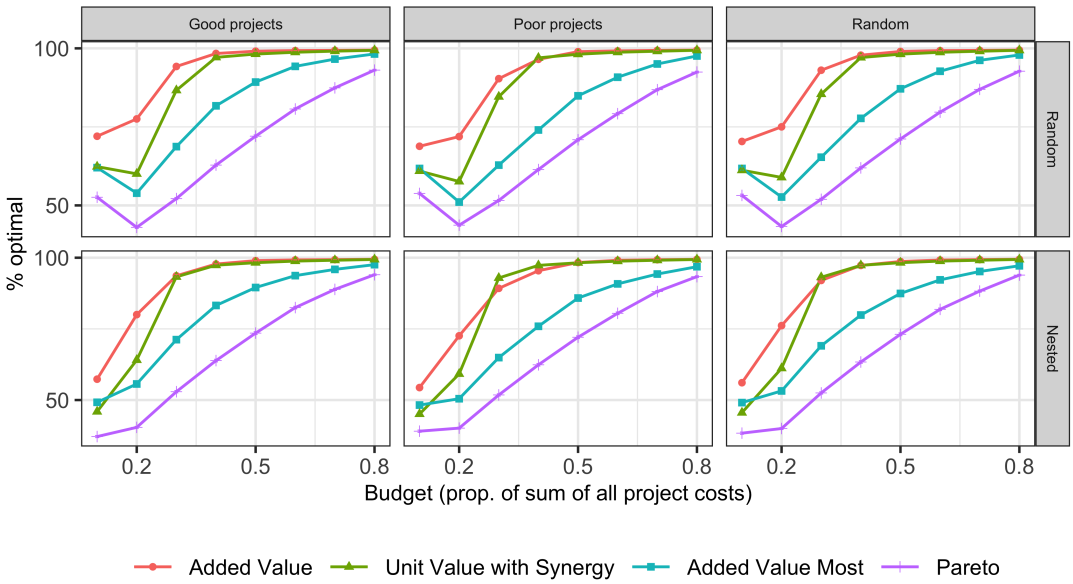

Figure 4: Mean relative portfolio value as a function of how projects interact with one another, for the best-performing fast and frugal portfolio heuristics. Plots in the bottom (top) row indicate whether higher-order interactions are nested within lower-order ones, or are random. Plots in different columns denote whether projects involved in interactions have high value-to-cost ratios (i.e. are "good" projects), low value-to-cost ratios ("poor" projects), or whether the selection is random.

This difference is erased and indeed reversed by the time budget levels reach 20% of total costs, with differences remaining small as budgets increase further. Thus the improvement in these two heuristics as budgets are initially increased from very low levels is larger when interactions are nested.

Both *Added Value* and *Added Value Most* perform better when interactions are constructed from “good” projects with high value-to-cost ratios than from relatively “poor” projects (Figure 4). Differences between “good” and “poor” interaction conditions are larger at lower budgets for the *Added Value* heuristic, but are relatively constant over budget conditions for *Added Value Most*. For both heuristics the random case occupies an intermediate condition between “good” and “poor”.

6 Behavioural study of portfolio decision making

6.1 Task description

We presented 75 participants with two versions of a simple portfolio selection task (the same one used in the numerical illustration in Appendix A). One version of the task was exactly the same as the example (Task 2); in the other version no project interactions were present (Task 1). Participants saw tasks in random order, were students from the African Institute of Mathematics and the University of the Western Cape, and were paid approximately $4 for their participation. Data collection errors occurred for two and one participants' in Task 1 and 2 respectively, leaving 73 and 74 participants respectively.

The task was worded generically, with no reference to any particular application area, to avoid biasing responses. Participants were instructed to choose a subset of “projects” that would collectively give them as many “points” as possible, subject to the same budget of 7 units. Participants were explicitly told that interactions existed between projects in some of the tasks, but were not told which projects were involved or the magnitude of the interactions – to do so would, in our opinion, bias responses and make the problem somewhat trivial. The decision problem thus involves an element of information gathering, because participants can only assess whether projects interact by selecting them, and in both tasks participants were allowed to remove or add projects. This has implications for analysis, which we discuss below.

Tasks were performed individually on a computer using an R Shiny web application (Chang et al., 2020). The interface consisted of a set of checkboxes in which participants could add or remove projects from their portfolios, and tables showing (a) individual project values and costs, (b) for each project not in the portfolio, the incremental change in portfolio value and cost that would result from its selection; (c) for each project in the portfolio, the incremental change in portfolio value and cost that would result from its deselection, (d) the current value and remaining budget of the currently selected portfolio. Part (a) is fixed but (b) – (d) depend on the current portfolio and are thus updated each time a project is selected or deselected. Each selection and deselection made by a respondent was recorded with an timestamp, and in this way it was possible to reconstruct the order in which projects were added or removed. When participants were satisfied with their chosen portfolio they clicked a button to submit their selection. The experimental interface was written in R 3.6.0 using `shiny` (Chang et al., 2020); results plots make use of packages `ggplot2` (Wickham, 2016) and `ggalluvial` (Brunson, 2020). All data and code used to set up the task and analyze responses are available at `https://github.com/iandurbach/portfolio-heuristics`.

6.2 Analysis

The assessment of the use of heuristics empirically faces problems of identifiability. The same project can be selected by different heuristics, and a random selection may lead to the same selection as any heuristic. Furthermore, because participants were not told which projects had interactions, some selections and deselections will be made with the purpose of gathering this information. In the absence of a search cost, it is not clear how much searching participants "should" do. We therefore analyzed both the final submitted portfolios as well as the order in which projects where added or removed before the final submission. For each respondent, we linked each project addition to a set of *potential heuristics* i.e. heuristics that would have selected the same project as was added, from the heuristics *Unit Value*, *Highest Value*, *Lowest Cost*, and *Added Value*. This association took into account the state of the current portfolio i.e. the projects already selected. Each project addition was allocated a single "vote"; in cases where the added project was selected by more than one heuristic, the vote was shared evenly between those heuristics. If the selection was not compatible with any heuristics it was allocated to an "other" category. Over all participants, this gave the weighted proportion of all selections that were consistent with the use of a particular heuristic. We excluded the *Unit Value with Synergy* and *Pareto* heuristics from this analysis as our collected data does not allow us to infer whether participants restricted their choices to interacting and non-dominated projects respectively.

We compared these proportions to what might be expected under a null model in which projects are added and removed at random. We did this by simulating a hypothetical sample of participants (of the same size as the real sample), with the same distribution of project additions and removals as observed in the experiment. For each participant, we added projects at random until the budget was exceeded. We then removed the project whose selection led to the budget violation, as well as one further project selected at random. We repeated this procedure of adding and removing projects until the desired number of removals had been achieved. The next time the budget was exceeded we removed the offending project and selected the remaining projects as the final portfolio. Once the hypothetical sample had been constructed in this way we calculated the proportion of selections consistent with each heuristic, in the same way as done for the true sample. We repeated this process 2000 times to create a distribution of proportions associated with each heuristic, under the null "random selection" model.

6.3 Results

The majority of participants' submitted portfolios that were consistent with portfolios selected by one of five major heuristics *Highest Value*, *Lowest Cost*, *Unit Value*, *Unit Value with Synergy*, or *Added Value* (Task 1: 55/73; Task 2: 61/74, see Table 1). In both tasks the most frequently selected portfolio consisted of $\{P_1, P_3, P_5\}$, which was selected by the *Unit Value* heuristic and was one of three possible portfolios selected by the *Highest Value* heuristic. The *Lowest Cost* and *Added Value* portfolios were rarely selected. In Task 1, 51/73 participants selected one of the optimal portfolios; in the more difficult Task 2 this proportion fell to 16/74. The sum of additions and removals, which can be considered a measure of participant effort, was positively associated with decision quality in both tasks but was particularly strong in Task 2, where participants selecting the optimal portfolio $\{P_1, P_2, P_3\}$ made on average 17.6 selections and deselections, compared to the sample mean of 7.7 (Table 1).

Of the 34 participants who chose portfolio $\{P_1, P_3, P_5\}$ in Task 2, the majority added projects in the same order as the *Highest Value* heuristic (5-3-1, 13/34 participants) or the *Unit Value* heuristic (5-1-3, 9/34 participants, see Table 2). Only 3 of the 16 participants who chose the optimal portfolio chose projects in the same order as predicted by *Unit Value with Synergy* (1-3-2), although no ordering was particularly popular. In Task 1 the most frequent ordering was not associated with any heuristic (1-3-5, 11/33 participants), with the second most frequent following the *Highest Value* heuristic (5-3-1, 10/33 heuristics). Other portfolios selected by the

Highest Value heuristic tended most often to have projects selected in the order dictated by the heuristic (Table 2).

$\mathbf{z}$	Heuristics supported	n	$V(\mathbf{z})$	$C(\mathbf{z})$	$\bar{s}_a$	$\bar{s}_r$
Task 1 (no interactions):						
135	uv,hv	33	8	6	4.4	1.4
235	hv	13	8	7	4.5	1.5
145	hv	5	8	7	6.2	3.4
124	–	5	4	7	3.8	0.8
125	lc	4	7	5	5.5	2.5
Task 2 (with interactions):						
135	uv,hv	34	11	6	4.4	1.4
123	sy	16	13	6	10.2	7.4
235	hv	8	8	7	3.8	0.8
34	–	4	4	7	2.0	0.0
125	av,lc	3	10	5	4.3	1.3

Table 1: Properties of the most frequently chosen portfolios in each task condition. For each portfolio $\mathbf{z}$ (shown using subscripts of selected projects) we show the number of participants choosing that portfolio, n, the set of heuristics that select $\mathbf{z}$ (hv = *Highest Value*, lc = *Lowest Cost*, uv = *Unit Value*, av = *Added Value*, sy = *Unit Value with Synergy*), portfolio value $V(\mathbf{z})$ and cost $C(\mathbf{z})$, and the mean number of selections (project additions) and deselections (removals) performed by participants during the experiment, $\bar{s}_a$ and $\bar{s}_r$, the sum of which can be considered a measure of effort. Optimal portfolios in each task are indicated in bold.

In both tasks the projects most frequently selected first were P_5 or P_1 (Task 1: P_5, 29/73; P_1, 25/73. Task 2: P_5, 35/73; P_1, 19/73, see Figure 5). Project P_5 is selected first by either *Highest Value* or *Unit Value* heuristics, while P_1 is selected by *Lowest Cost.* Regardless of which project was selected first the project most commonly added next was P_3, which in Task 1 is the project selected by *Unit Value* and one of two projects selected by *Highest Value.* In Task 2 P_3 is also selected by *Added Value* if P_1 is selected first (Task 1: 27/29; Task 2: 33/35). Subsequent additions are much more evenly distributed over projects as the choice becomes more heavily influenced by which projects are already in the portfolio. The most common initial additions are 1-3-5, 5-3-1 and 5-3-2 in Task 1 (10, 7 and 6 participants respectively, see Figure 5), and 5-3-1, 5-1-3 and 1-3-5 in Task 2 (16, 8, and 8 participants respectively). As mentioned, 5-3-1 and 5-3-2 are both consistent with the *Highest Value* heuristic, while 5-1-3 is consistent with *Unit Value.*

The proportion of selections that were consistent with the *Highest Value* or *Unit Value* heuristics in Task 1, and with the *Unit Value*, *Added Value*, and *Highest Value* heuristics in Task 2, are very unlikely to arise from a random selection strategy (Task 1: $p = 1/2000$ and $p < 1/2000$ respectively; Task 2: $p = 3/2000$, $p = 10/2000$, $p = 113/2000$ respectively, see Figure

	Task 1: no interactions				Task 2: with interactions		
z	Order R1	Order R2	Order R3	**z**	Order R1	Order R2	Order R3
135	1-3-5 (11)	5-3-1 (10)	3-5-1 (5)	135	5-3-1 (13)	5-1-3 (9)	1-3-5 (7)
235	5-3-2 (7)	5-2-3 (3)	2-3-5 (2)	123	3-1-2 (5)	1-2-3 (4)	1-3-2 (3)
145	5-4-1 (2)	4-5-1 (2)	5-1-4 (1)	235	5-3-2 (5)	3-5-2 (1)	2-3-5 (1)
124	1-2-4 (3)	2-1-4 (1)	4-1-2 (1)	34	4-3 (3)	3-4 (1)	
125	5-2-1 (1)	1-5-2 (1)	2-5-1 (1)	125	2-5-1 (1)	1-2-5 (1)	5-1-2 (1)

Table 2: Selection order for projects appearing in the most frequently chosen portfolios. For each portfolio **z** we show the order in which the projects making up the portofolio were added. We show the three most popular orderings, which in most cases account for the majority of participants. The number of participants using each sequence is shown in parentheses.

6). Similarly, a much lower proportion of selections could not be explained by any heuristics than would be expected if selections were made randomly ($p < 1/2000$, see the "Other" column of Figure 6). While variation from a random strategy is not a particularly stringent hurdle, in conjunction with our other results these provide some evidence that unassisted decision makers are employing at least some of the heuristics we propose in this study. We also examined consecutive selections and assessed the proportion of opportunities to complete an interaction subset that were taken. Participants were more likely to select a project that completed one of the two-project interactions i.e. 1-2, 1-3, in Task 2 than in Task 1, suggesting that interaction information was used (Task 1: 61/121 selections (50%), Task 2: 98/156 selections (63%), $z = 2.1$, $p = 0.04$). This proportion increased further to 73% (42/58) if the project was also the *Added Value* selection.

7 Conclusions and further research

Portfolio decisions are an important and increasingly studied class of decision problem, with optimization models developed for a variety of settings (e.g. Salo et al., 2011; Cranmer et al., 2018; Vilkkumaa et al., 2018). We see two gaps in this literature. Firstly, portfolio optimization typically means that one has to assess all project interactions. The effort involved in this can be considerable and, even in a prescriptive setting, it is reasonable that decision makers might want to limit this. There is currently relatively little guidance from portfolio decision analysis for how to do so. Secondly, relatively little is known about how people actually go about making portfolio decisions involving project interactions (Fasolo et al., 2011; Phillips and Bana e Costa, 2007; Schiffels et al., 2018).

Heuristics have played an important role in addressing these two issues in conventional

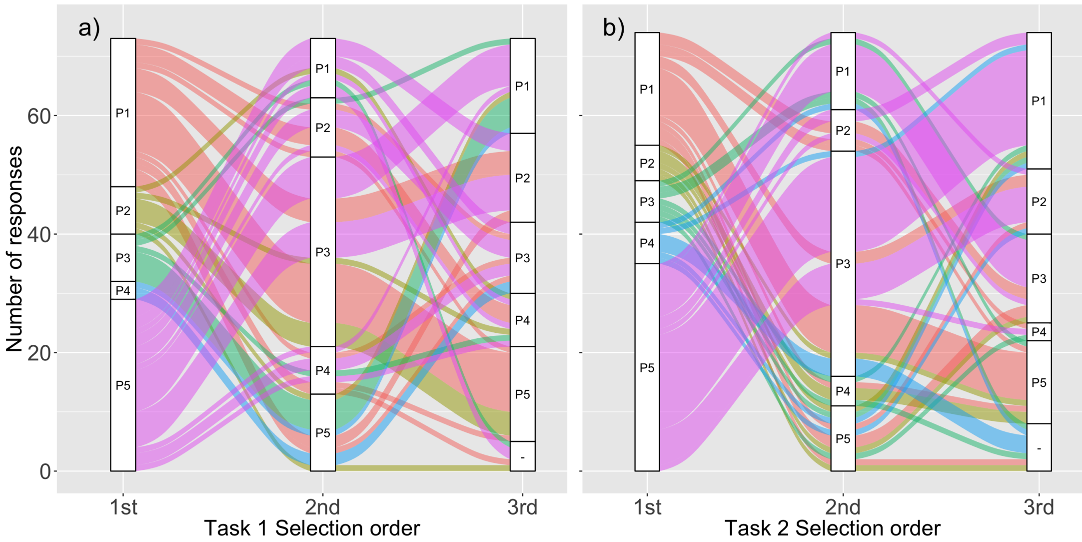

Figure 5: Visualizing the frequencies of the first three selections made. The height of a block represents the number of participants who selected that project in a particular position (1st, 2nd, 3rd). The width of a stream between two projects represents the number of participants who chose both projects in the respective positions traversed by the stream. The colour of a stream denotes the first project chosen.

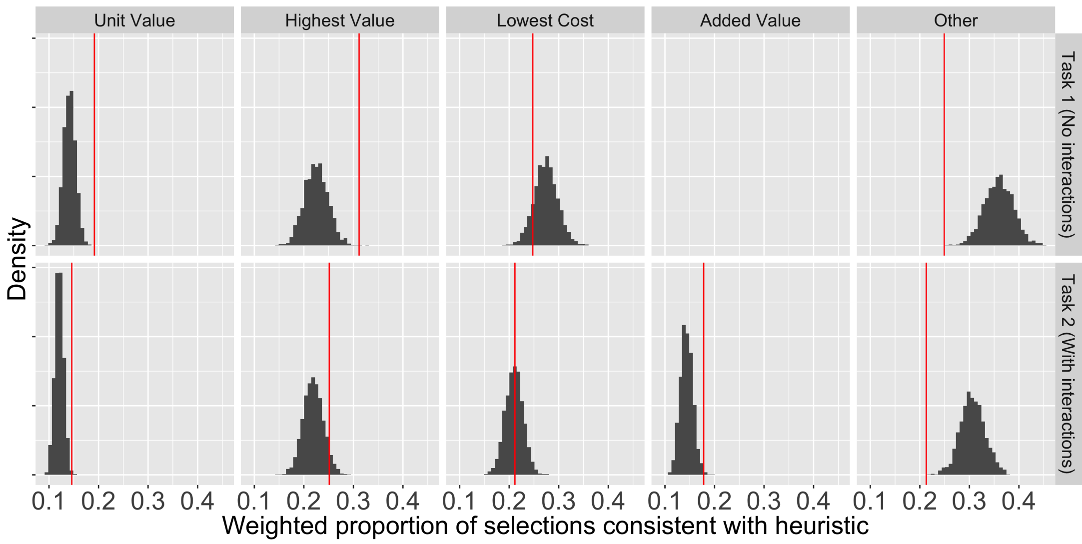

Figure 6: Proportion of project selections that were consistent with each heuristic (red vertical lines). As at any stage in the process different heuristics can select the same project, these proportions are of limited value on their own. We therefore compare each one against a distribution of proportions generated by a random selection heuristic (grey histograms; see text for details). In cases where the same project is selected by different heuristics, that selection's "vote" is distributed evenly between those heuristics, and hence the proportion is a weighted one

one-out-of-n decisions (e.g. Tversky and Kahneman, 1974; Hogarth and Karelaia, 2005, 2006), and there is every reason to think that they may be useful for portfolio decision making too. Ours is not the first paper to study portfolio heuristics (Keisler, 2004, 2005, 2008; Schiffels et al., 2018), but we do propose a number of new heuristics, include the key issue of project interactions, and use a multi-method approach employing simulation, analytical results, and behavioral experiment. This provides a more detailed understanding of the potential benefits of heuristics in finding a balance between the effort required to assess all possible interactions and the value of the selected portfolio.

Analytical results showed that heuristics require a small fraction of the assessments needed for exact methods. Nevertheless, the number of assessments can still be large, at least for the *Added Value* heuristic at most realistic problem settings. This is indicative of the complexity of portfolio decision making, and the poor performance of heuristics that ignore interactions show the price to be paid for more extreme frugality. Still, it is not entirely clear how "fast" the *Added Value* heuristic could be, if for example interactions must be constantly evaluated but are time-consuming to assess. The *Unit Value with Synergy* heuristic would appear to be more frugal and thus to offer a more intuitively attractive balance between assessment effort and portfolio value, although it is difficult to precisely specify its information requirements. The heuristic of course depends strongly on interactions between projects being positive. How best to incorporate negative and other forms of project interactions is a topic we leave to future research.

Our simulation results showed that two heuristics, *Added Value* and *Unit Value with Synergy* provided outcomes that were competitive with theoretically optimal models under a fairly wide range of environmental conditions. Conclusions drawn from our simulations are, as with all simulations, heavily dependent on the ranges of assumed parameter values, but provide initial evidence that at least these two heuristics may provide trade-offs between assessment effort and portfolio value that could be viewed favourably by decision makers. The two heuristics performed best when interactions between projects were nested rather than random (except at very low budgets), and when positive interactions existed primarily between projects that were also individually good. These specify the conditions under which it would be ecologically rational (Gigerenzer et al., 1999) to use either heuristic and thus features that a future empirical study of real-world portfolio decisions might search for. The mostly extremely poor performance of all heuristics ignoring interactions, including the *Pareto* heuristic, is an important and somewhat surprising negative result.

Studying portfolio decision making in a laboratory context is difficult because the experimenter is faced with a choice between making all project interactions known (in which case the key issue of interaction assessment is ignored, and responses likely biased) or not (in which case responses are a mixture of gathering information on interactions and statements of preference). Our choice was the latter, and we assessed results by examining the final portfolios selected and by comparing project additions to what would be expected under a random selection strategy. Our results showed that (a) participants tended to choose certain portfolios more often than would be expected by chance alone, and that these portfolios were the same as those selected by our *Unit Value* or *Highest Value* heuristics, (b) a greater-than-chance proportion of participants who chose these portfolios added the projects making up the portfolios in the same order as the two heuristics, and (c) the most popular initial selections of projects were also consistent with *Unit Value* or *Highest Value* heuristics. Our findings are in broad agreement with what Schiffels et al. (2018) found for portfolio problems without interactions – we also find common use of *Unit Value* (although not *Lowest Cost*) and substantial variability of heuristic use both between and within participants.

Our core result is that psychologically plausible heuristics can select excellent portfolios using a fraction of the information required by optimal methods, but they must use at least some interaction information to do so. Crucially, it appears that a little interaction information goes a long way; in our simulated contexts it was more important to know which projects were involved in *any* positive interaction than to estimate the magnitude of those interactions. Our work suggests two possible modes for using portfolio heuristics in the broader context of a portfolio decision support system (Ghasemzadeh and Archer, 2000; Lourenco et al., 2012; Jang, 2019; Kreuzer et al., 2020). The first mode views portfolio heuristics as a drop-in replacement for more information-intensive optimization methods, appropriate for applications where time or other constraints make it impossible to assess the information required by optimization methods. Portfolio heuristics are computationally straightforward to implement and decision support facilitating the application of a particular heuristic follows more-or-less directly from the heuristic's definition. Implementation of *Unit Value with Synergy* requires an initial step in which the set of candidate projects is pruned to include only those projects with any positive interactions, followed by a second step establishing the value-to-cost ratios of those projects, following which projects are added greedily. Implementation of *Added Value* requires the initial assessment of individual projects' values and costs, and ranking by their value-to-cost ratios. After each addition of a project to the portfolio, an assessment round is required to collect data on any interactions between the project just included and the remaining candidate projects, after which value-to-cost ratios of candidate projects can be updated and the next addition made. The second mode is to use portfolios selected by fast and frugal heuristics as a basis for comparison with portfolios selected by exact methods, where all interaction information is available. Decision support systems for portfolio decision making routinely include value-to-cost ratios, and include a comparison with portfolios constructed on a greedy basis from these data (e.g. PROBE, Lourenco et al., 2012). Fast and frugal heuristics augment these sources of comparative information and also allow one to estimate the value of assessing interaction information beyond that required by portfolio heuristics, in the manner of Keisler (2004, 2008).

Our study suggests a number of promising avenues for further work: characterizing the features of real-world portfolio decisions, incorporating other kinds of interactions between projects, incorporating multiple attributes and uncertainties, and developing assessment procedures for *Unit Value with Synergy*. Given our results on the importance of project interactions, development of further heuristics is probably best aimed at heuristics that simplify interaction information in some way. Most of the heuristics considered in this paper are single-cue heuristics that use one piece of information to discriminate between options, but the good performance offered by our one multiple cue heuristic (*Unit Value with Synergy*, which lexicographically considers the potential for positive interaction and unit value) suggests that combining cues in imaginative ways may be a fruitful way to reduce information requirements.

Acknowledgements

ID is supported in part by funding from the National Research Foundation of South Africa (Grant ID 90782, 105782).

References

M. Airoldi and A. Morton. Portfolio decision analysis for population health. In A. Salo, J. Keisler, and A. Morton, editors, *Portfolio Decision Analysis*, pages 359–381. Springer, 2011.

N. Arratia, F. López, S. Schaeffer, and L. Cruz-Reyes. Static R&D project portfolio selection in public organizations. *Decision Support Systems*, 84:53–63, 2016.

E. Baker, V. Bosetti, and A. Salo. Robust portfolio decision analysis: An application to the energy research and development portfolio problem. *European Journal of Operational Research*, 284:1107–1120, 2020.

C. T. Ball, H. J. Langholtz, J. Auble, and B. Sopchak. Resource-allocation strategies: A verbal protocol analysis. *Organizational Behavior and Human Decision Processes*, 76(1):70–88, 1998.

M. Barbati, S. Greco, M. Kadziński, and R. Słowiński. Optimization of multiple satisfaction levels in portfolio decision analysis. *Omega*, 78:192–204, 2018.

M. Baucells, J. A. Carrasco, and R. M. Hogarth. Cumulative dominance and heuristic performance in binary multiattribute choice. *Operations Research*, 56(5):1289–1304, 2008.

A. Bröder and B. Newell. Challenging some common beliefs: Empirical work within the adaptive toolbox metaphor. *Judgment and Decision Making*, 3(3):205214, 2008.

J. C. Brunson. ggalluvial: Layered grammar for alluvial plots. *Journal of Open Source Software*, 5(49):2017, 2020.

M. Buckmann and Ö. Şimşek. Decision heuristics for comparison: How good are they? In T. V. Guy, M. Krn, D. Rios-Insua, and D. H. Wolpert, editors, *Proceedings of the NIPS 2016 Workshop on Imperfect Decision Makers*, volume 58, pages 1–11. PMLR, 2017.

W. Chang, J. Cheng, J. Allaire, Y. Xie, and J. McPherson. *shiny: Web Application Framework for R*, 2020. URL `https://CRAN.R-project.org/package=shiny`. R package version 1.4.0.2.

A. Cranmer, E. Baker, J. Liesiö, and A. Salo. A portfolio model for siting offshore wind farms with economic and environmental objectives. *European Journal of Operational Research*, 267 (1):304314, 2018.

B. Fasolo, A. Morton, and D. von Winterfeldt. Behavioural issues in portfolio decision analysis. In A. Salo, J. Keisler, and A. Morton, editors, *Portfolio Decision Analysis*, pages 149–165. Springer, 2011.

J. K. Ford, N. Schmitt, S. L. Schechtman, B. M. Hults, and M. L. Doherty. Process tracing methods: Contributions, problems, and neglected research questions. *Organizational Behavior and Human Decision Processes*, 43(1):75–117, 1989.

F. Ghasemzadeh and N. P. Archer. Project portfolio selection through decision support. *Decision support systems*, 29(1):73–88, 2000.

G. Gigerenzer and D. G. Goldstein. Reasoning the fast and frugal way: models of bounded rationality. *Psychological review*, 103(4):650, 1996.

G. Gigerenzer, P. M. Todd, and A. R. Group. *Simple heuristics that make us smart.* Oxford University Press, 1999.

F. Hassanzadeh, H. Nemati, and M. Sun. Robust optimization for interactive multiobjective programming with imprecise information applied to r&d project portfolio selection. *European Journal of Operational Research*, 238(1):41–53, 2014.

R. M. Hogarth and N. Karelaia. Simple models for multiattribute choice with many alternatives: When it does and does not pay to face trade-offs with binary attributes. *Management Science*, 51(12):1860–1872, 2005.

R. M. Hogarth and N. Karelaia. Regions of rationality: Maps for bounded agents. *Decision Analysis*, 3(3):124–144, 2006.

H. Jang. A decision support framework for robust r&d budget allocation using machine learning and optimization. *Decision Support Systems*, 121:1–12, 2019.

U. Jung and D. Seo. An anp approach for r&d project evaluation based on interdependencies between research objectives and evaluation criteria. *Decision Support Systems*, 49(3):335–342, 2010.

K. V. Katsikopoulos. Psychological heuristics for making inferences: Definition, performance, and the emerging theory and practice. *Decision Analysis*, 8(1):10–29, 2011.

K. V. Katsikopoulos, I. N. Durbach, and T. J. Stewart. When should we use simple decision models? a synthesis of various research strands. *Omega*, 81:17–25, 2018.

J. Keisler. Value of information in portfolio decision analysis. *Decision analysis*, 1(3):177–189, 2004.

J. Keisler. When to consider synergies in project portfolio decisions. College of Management Working Papers and Reports 11, University of Massachusetts Boston, 2005. URL `http://scholarworks.umb.edu/management_wp/11`.

J. M. Keisler. The value of assessing weights in multi-criteria portfolio decision analysis. *Journal of Multi-Criteria Decision Analysis*, 15(5-6):111–123, 2008.

N. Keller and K. V. Katsikopoulos. On the role of psychological heuristics in operational research; and a demonstration in military stability operations. *European Journal of Operational Research*, 249(3):1063–1073, 2016.

D. N. Kleinmuntz. Resource allocation decisions. In W. Edwards, R. Miles, and D. von Winterfeldt, editors, *Advances in decision analysis: From foundations to applications*, chapter 20, pages 400–418. Cambridge University Press, 2007.

V. Korotkov and D. Wu. Evaluating the quality of solutions in project portfolio selection. *Omega*, 91:102029, 2020.

T. Kreuzer, M. Röglinger, and L. Rupprecht. Customer-centric prioritization of process improvement projects. *Decision Support Systems*, pages 1–14, 2020. URL `https://doi.org/10.1016/j.dss.2020.113286`.

H. Langholtz, C. Gettys, and B. Foote. Resource-allocation behavior under certainty, risk, and uncertainty. *Organizational behavior and human decision processes*, 54(2):203–224, 1993.

H. J. Langholtz, C. Ball, B. Sopchak, and J. Auble. Resource-allocation behavior in complex but commonplace tasks. *Organizational Behavior and Human Decision Processes*, 70(3):249–266, 1997.

H. J. Langholtz, A. T. Marty, C. T. Ball, and E. C. Nolan. *Resource-allocation behavior*. Springer Science & Business Media, 2002.

J. Liesiö. Measurable multiattribute value functions for portfolio decision analysis. *Decision Analysis*, 11(1):1–20, 2014.

J. Liesiö and A. Salo. Scenario-based portfolio selection of investment projects with incomplete probability and utility information. *European Journal of Operational Research*, 217(1):162–172, 2012.

J. Liesiö, P. Mild, and A. Salo. Preference programming for robust portfolio modeling and project selection. *European Journal of Operational Research*, 181(3):1488–1505, 2007.

J. C. Lourenco, A. Morton, and C. A. B. e Costa. Probea multicriteria decision support system for portfolio robustness evaluation. *Decision support systems*, 54(1):534–550, 2012.

P. Mild, J. Liesiö, and A. Salo. Selecting infrastructure maintenance projects with robust portfolio modeling. *Decision Support Systems*, 77:21–30, 2015.

T. Pape, S. Kavadias, and S. C. Sommer. The Behavioural Knapsack Problem: Evidence of a Fundamental Bias in Project Selection. 2019. URL `http://dx.doi.org/10.2139/ssrn.3448676`.

L. Phillips and C. Bana e Costa. Transparent prioritisation, budgeting and resource allocation with multi-criteria decision analysis and decision conferencing. *Annals of Operations Research*, 154(1):51–68, 2007.

D. Pisinger. The quadratic knapsack problema survey. *Discrete applied mathematics*, 155(5): 623–648, 2007.

J. Rieskamp, J. R. Busemeyer, and T. Laine. How do people learn to allocate resources? comparing two learning theories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(6):1066, 2003.

A. Salo, J. Keisler, and A. Morton. An invitation to portfolio decision analysis. In A. Salo, J. Keisler, and A. Morton, editors, *Portfolio Decision Analysis*, pages 3–27. Springer, 2011.

S. Schiffels, T. Fliedner, and R. Kolisch. Human behavior in project portfolio selection: Insights from an experimental study. *Decision Sciences*, 49(6):1061–1087, 2018.

Ö. Şimşek. Linear decision rule as aspiration for simple decision heuristics. In *Advances in Neural Information Processing Systems*, pages 2904–2912, 2013.

C. Stummer and K. Heidenberger. Interactive R&D portfolio analysis with project interdependencies and time profiles of multiple objectives. *IEEE Transactions on Engineering Management*, 50(2):175–183, 2003.

S. Theussl and K. Hornik. *Rglpk: R/GNU Linear Programming Kit Interface*, 2019. URL `https://CRAN.R-project.org/package=Rglpk`. R package version 0.6-4.

P. M. Todd. How much information do we need? *European Journal of Operational Research*, 177(3):1317–1332, 2007.

A. Tversky. Elimination by aspects: A theory of choice. *Psychological review*, 79(4):281–299, 1972.

A. Tversky and D. Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, 1974.

E. Vilkkumaa, J. Liesiö, A. Salo, and L. Ilmola-Sheppard. Scenario-based portfolio model for building robust and proactive strategies. *European Journal of Operational Research*, 266(1): 205–220, 2018.

H. Wickham. *ggplot2: Elegant Graphics for Data Analysis.* Springer-Verlag New York, 2016. ISBN 978-3-319-24277-4. URL `https://ggplot2.tidyverse.org`.

A Numerical illustration of add-the-best heuristics

Suppose that a decision maker must construct a portfolio from five projects P_1–P_5 with values and costs given in Table A.1. Positive interactions exist between the following subsets of projects: P_1, P_2, P_3 (interaction subset $\mathcal{A}_1$); P_2, P_3, P_4 (interaction subset $\mathcal{A}_2$); P_1, P_2 (interaction subset $\mathcal{A}_3$); P_1, P_3 (interaction subset $\mathcal{A}_4$). If all of the projects in any of these interaction subsets are selected, an additional value of $B = 3$ is added to the value of the portfolio. The decision maker has a budget of $\zeta = 7$. The optimal solution is to select P_1, P_2, P_3, which returns a portfolio value of 13 at a cost of 6.

			Unit Value				*Highest Value*				*Lowest Cost*			
			Criterion value at stage				Criterion value at stage				Criterion value at stage			
	b_j	c_j	0	1	2	3	0	1	2	3	0	1	2	3
P_1	1	1	1/1	1/1	–	–	1	1	1	–	1	–	–	–
P_2	1	2	1/2	1/2	1/2	1/2*	1	1	1*	1*	2	2	–	–
P_3	2	3	2/3	2/3	2/3	–	2	2	2*	2*	3	3	3	3*
P_4	2	4	1/2	1/2	1/2	1/2*	2	2	–	–	4	4	4	4*
P_5	5	2	5/2	–	–	–	5	–	–	–	2	2	2	–
Selection			P_5	P_1	P_3	–	P_5	P_4	P_1	–	P_1	P_2	P_5	–

Table A.1: A numerical illustration of proposed fast and frugal portfolio heuristics ignoring quantitative interaction information. Relevant columns show the information required by each heuristic at each iteration i.e. as projects are sequentially added to the portfolio (project values, costs, and the ratio between the two for *Highest Value*, *Lowest Cost*, and *Unit Value* respectively). Projects that cannot be added due to budget constraints are indicated with an asterisk.

The *Highest Value* heuristic selects projects in decreasing order of value. In our example it first adds P_5 and then picks randomly between P_4 and P_3. If P_4 is chosen only P_1 can be chosen without exceeding the budget. If P_3 is chosen after P_5 then two units of budget remain and either P_1 or P_2 (which have the same value) can be chosen. Thus *Highest Value* can select any of the portfolios $\{P_5, P_4, P_1\}$, $\{P_5, P_3, P_2\}$, or $\{P_5, P_3, P_1\}$, which have values 8, 8, and 11 and costs 7, 7, and 6, respectively.

The *Lowest Cost* heuristic starts by selecting the cheapest project, P_1. The next cheapest projects, P_2 and P_5, both have a cost of two and are thus added in either order. Adding any other project would exceed the budget so the final selection is $\{P_1, P_2, P_5\}$, which has a value of 10 and a cost of 5.

The *Unit Value* heuristic sequentially adds projects P_5, P_1, and P_3, after which the cost of both remaining projects exceeds the available budget. The selected portfolio has a total value of 11 (8 for the value of each of the projects plus the value of interaction $\mathcal{A}_4$) and a cost of 6.

The *Pareto* heuristic involves a random selection from the set of non-dominated candidates at each step. Suppose the first candidate is P_2. As it is dominated by P_1, P_2 is not chosen and a new candidate it randomly chosen. Suppose that P_1 is now picked; it is non-dominated and thus selected. Suppose that P_2 is again randomly selected as the next candidate. Although P_2 is dominated by P_1, P_1 is already in the portfolio and thus, because it is not dominated by any other candidate and is within budget, P_2 would be selected. After selecting P_2, P_4 could not be accepted because it is dominated by P_3 but P_3 and P_5 are equally likely to be selected in the next and final step. These portfolios have values of 13 and 10 and costs of 6 and 5, respectively.

The *Unit Value with Synergy* heuristic first identifies any project that has a positive interaction with another project – all projects except for P_5. It then adds projects in this set using the *Unit Value* heuristic, that is by their individual value-to-cost ratios, and thus adds P_1, P_3, and P_2 (since P_4 would exceed the available budget). The selected portfolio is the optimal one.

The *Added Value* heuristic first adds P_5 and P_1, which give the biggest increases in portfolio value per unit cost (there are no two-project interactions). After this there are two interaction subsets that may be completed by the addition of a new project: interaction subset $\mathcal{A}_3$ would be completed by adding P_2 while interaction subset $\mathcal{A}_4$ would be completed by adding P_3. Adding P_2 increases portfolio value by 4 at a cost of 2 while adding P_3 increases value by 5 at a cost of 3 (Table A.2). Thus P_2 is selected. Adding any other candidate project would exceed the available budget of 7 and so the final selection is $\{P_5, P_1, P_2\}$, giving a value of 10 at a cost of 5.

Added Value Most, *Added Value Least*, and *Added Value Random* all begin by adding P_5 and then, as P_5 does not belong to any interaction subsets, P_1. The three then diverge. *Added Value*

			Added Value				*Added Value Most*				*Added Value Least*			
			Criterion value at stage				Criterion value at stage				Criterion value at stage			
Proj	b_j	c_j	0	1	2	3	0	1	2	3	0	1	2	3
P_1	1	1	1	1	–	–	1	1	–	–	1	1	–	-
P_2	1	2	1/2	1/2	2/1	–	1/2	1/2	1/2	1/2*	1/2	1/2	2/1	-
P_3	2	3	2/3	2/3	5/3	8/3*	2/3	2/3	2/3	–	2/3	2/3	5/3	8/3*
P_4	2	4	1/2	1/2	1/2	1/2*	1/2	1/2	1/2	1/2*	1/2	1/2	1/2	1/2*
P_5	5	2	5/2	–	–	–	5/2	–	–	–	5/2	–	-	–
Selection			P_5	P_1	P_2	–	P_5	P_1	P_3	–	P_5	P_1	P_2	–

Table A.2: A numerical illustration of proposed fast and frugal portfolio heuristics making use of quantitative interaction information. The table shows, at each decision stage, the criterion value assigned by each heuristic to each of the eligible projects (i.e. the estimated increase in portfolio value per unit cost as projects are sequentially added to the portfolio). Projects that cannot be added due to budget constraints are indicated with a superscripted asterisk.

Most identifies the most valuable of the already included projects, which is P_5. It therefore does not need to update the values of the remaining projects, since P_5 has no possible interactions with any of them (see Table A.2). Thus the next project added is P_3. Further selections exceed the budget, and the selected portfolio $\{P_5, P_1, P_3\}$ has a value of 11 and a cost of 6.

Added Value Least considers only the interactions involving the least valuable project in the portfolio (P_1). This makes project P_2 and P_3 more attractive because of the completable interaction sets $\mathcal{A}_3 = \{P_1, P_2\}$ and $\mathcal{A}_4 = \{P_1, P_3\}$. Project P_2 is selected next, after which no further projects are within budget. The final selection is $\{P_5, P_1, P_2\}$, giving a value of 10 at a cost of 5. Updates to the value-cost ratios are shown in Table A.2.

Added Value Random randomly chooses one of them: only interactions with the selected project will be considered in the next step. If P_5 is chosen then the heuristic selects P_3 next. It then randomly chooses between P_5, P_3, and P_1, again only considering interactions with the selected project in the following step. Regardless of this choice, further selections exceed the budget, and the selected portfolio is $\{P_5, P_1, P_3\}$. If P_1 is randomly chosen in the first step then P_2 is added at the next step and the heuristic terminates.