Ex-Post Optimal Knapsack Procurement

Felix Jarman (German Ministry of Finance)
Vincent Meisner (TU Berlin)

Discussion Paper No. 27

March 29, 2017

Collaborative Research Center Transregio 190 | www.rationality-and-competition.de
Ludwig-Maximilians-Universität München | Humboldt-Universität zu Berlin
Spokesperson: Prof. Dr. Klaus M. Schmidt, University of Munich, 80539 Munich, Germany
+49 (89) 2180 3405 | info@rationality-and-competition.de

Ex-post Optimal Knapsack Procurement*

Felix Jarman[†] Vincent Meisner[‡]

March 28, 2017

Abstract

We consider a budget-constrained mechanism designer who selects an optimal set of projects to maximize her utility. Projects may differ in their value for the designer, and their cost is private information. In this allocation problem, the quantity of procured projects is endogenously determined by the mechanism. The designer faces ex-post constraints: The participation and budget constraints must hold for each possible outcome, while the mechanism must be strategyproof. We identify settings in which the class of optimal mechanisms has a deferred acceptance auction representation which allows an implementation with a descending-clock auction. Only in the case of symmetric projects do price clocks descend synchronously such that the cheapest projects are implemented. The case in which values or costs are asymmetrically distributed features a novel tradeoff between quantity and quality. The reason is that guaranteeing allocation to the most favorable projects under strategyproofness comes at the cost of a diminished expected number of conducted projects.

JEL-Classification: D02, D44, D45, D82, H57.
Keywords: Mechanism Design, Knapsack, Budget, Procurement, Auction, Deferred Acceptance Auctions

*We thank Ludwig Ensthaler, Daniel García, Thomas Giebe, Hans Peter Grüner, Doh-Shin Jeon, Volker Nocke, Ilya Segal, Christoph Siemroth, André Stenzel, Thomas Tröger, anonymous referees, and participants of conferences and seminars in Berlin, Bonn, Bielefeld, Lisbon, London, Louvain-la-Neuve, Mannheim, Oslo, Toulouse and Vienna. Support by German Science Foundation (DFG) through CRC 884 (Felix Jarman) and CRC TRR 190 (Vincent Meisner) is gratefully acknowledged.

[†]felixjarman@gmail.com

[‡]Technical University Berlin, Straße des 17. Juni 135, 10623 Berlin, Germany, vincent.meisner@tu-berlin.de

1 Introduction

We study the problem of a procurer who can spend a fixed budget on any of n available projects which differ in the value the procurer derives from them. Projects (agents) have private information about their costs and want to get funding beyond the necessary minimum. The designer's goal is to select an affordable set of maximal aggregate quality. In other words, she faces a mechanism design variant of the knapsack problem[1] with strategic behavior due to informational asymmetries. Essentially, we approach this problem as an "up to possibly n-units" procurement problem with n agents with single-unit supply where demand quantity is determined after observing projects' reports under a budget constraint. The budget constraint, the individual-rationality constraints, and the incentive-compatibility constraints are imposed ex-post, i.e., for any cost realization, the sum of transfers must not exceed the budget, implemented projects always have to be at least fully compensated, and truth-telling must be a (weakly) dominant strategy. We fully characterize the optimal mechanism in the two-project case and the symmetric case, and suggest an implementation with a descending-clock auction with a deferred acceptance rule. Furthermore, we discuss which insights of the two-project case carry over to the general asymmetric case and identify weakening the substitutes condition as the natural next step of this research endeavor. Because of a tradeoff between quantity and quality, an optimal price clock may have to stop for a period of time leading to instances in which an inferior project is implemented instead of a superior one.

This framework matches a large range of allocation problems, in which a designer needs to allocate a divisible but fixed capacity among agents. Allocation problems, in which a financial budget constraint represents the fixed capacity, include the procurement of bus lines, bridges, and streets, or the allocation of subsidies or research money. Alternatively, the capacity constraint can represent the payload limit on a freighter or on a space shuttle,[2] or a limited amount of time to be devoted to several tasks. Out of many suitable applications, we employ as our leading example a development fund that desires to distribute money to nonprofit projects with nonmonetary benefits.

[1] The knapsack problem is a classical combinatorial problem, dating as far back as 1897. A set of items is assigned values and weights. The knapsack should be filled with the maximal value, but can carry only up to a given weight. For an overview of the literature on knapsack problems, see Kellerer, Pferschy, and Pisinger (2004).

[2] Clearly, the capacity of a space shuttle is limited. The problem of optimally allocating the capacity and incentivizing projects to reduce payload is economically relevant, see Ledyard, Porter, and Wessen (2000).

Our paper not only helps to understand a class of economically relevant problems, the framework also presents a novel methodological challenge. The ex-post nature of both the participation and the budget constraint precludes the use of standard pointwise optimization techniques à la Myerson (1981). Nonetheless, rewriting the problem involves expressing expected transfers in terms of the allocation function as an auxiliary step. As the designer maximizes expected payoff including residual money, we can employ a procurement analogue of Myerson's notion of "virtual values". However, our results qualitatively translate to a setting in which the designer does not value residual money.

By focusing on strategyproof deterministic mechanisms, we can reduce the problem to finding a set of optimal cutoff functions z_i that, for each project i, map the cost vector of other projects $\mathbf{c}_{-i}$ into a cutoff cost level. Project i is conducted if and only if i's cost report falls weakly below cutoff $z_i(\mathbf{c}_{-i})$ and the corresponding compensation payment for that case equals this cutoff. Next, we investigate properties these cutoff functions exhibit in optimum. In the two-project case, the optimal allocation rule has substitutes: Given a project is implemented for some cost vector, it is also implemented when, all else being equal, the cost of the rival project is increased. This property may cease to be optimal in more general cases. However, we show that any optimal allocation rule with substitutes has non-bossy winners: A single project that is implemented cannot affect the allocation without changing its own allocation status. Finally, the optimal allocation rule excludes all projects with negative "virtual surplus" from the allocation.

By virtue of these properties, such a mechanism has an equivalent deferred acceptance (DA) auction representation as analyzed in Milgrom and Segal (2015). We show that project substitutability is not necessary and discuss a weaker form of substitutes. A DA auction is an iterative algorithm that computes the allocation and transfers of an auction mechanism and possesses attractive features with respect to bidders' incentives that go beyond dominant-strategy implementability. First, in any DA auction, revealing the type truthfully is an "obviously dominant strategy" as defined by Li (2015).[3] Second, any DA auction is weakly group-strategyproof. In other words, it is impossible for a coalition of projects to coordinate their bidding strategies such that it strictly increases the utility of all projects in the coalition. Third, the dominant-strategy equilibrium outcome of any DA auction is the only outcome that survives iterated deletion of dominated

[3] There does not exist any deviation such that, in any information set in which a deviating action is played, the best-case deviation payoff (against even the most favorable profile of strategies of the other players that is consistent with this information set) is strictly larger than the worst-case payoff from truthful bidding (achieved against the least favorable such strategy profile).

strategies in the corresponding full information game with the same allocation rule but where players pay their own bid. Therefore predicting the dominant-strategy equilibrium outcome in a DA auction can be considered robust.

Milgrom and Segal (2015) argue that these properties make DA auctions suitable for many challenging environments such as radio spectrum reallocations. Most importantly, they show that every DA auction can be represented by a descending-clock auction. Among several potential applications, they also consider our budget-constrained procurement setup (Example 7: "Budget Constraint"). However, they do not show optimality of the DA auction. Therefore we can strengthen the argument in favor of DA auctions. The techniques established in our paper may be helpful to prove optimality of DA auctions in the other settings mentioned in their paper.

The existence of a corresponding (direct) DA auction implies that the allocation rule can be implemented with an appropriately designed descending-clock auction as its corresponding indirect form: Every project faces a clock with a continuously descending price on it, and indicates whether it is willing to conduct its project at this price. Prices do not ascend again. In this auction, it is a weakly dominant strategy for any project to exit the auction once the clock price hits the project's actual cost level. We show that it is optimal to rank projects according to their cost and "greenlight" the cheapest ones, when projects have identical values and costs are drawn from the same distribution. That is, price clocks run down synchronously and hence projects exit in order of their costs until the budget suffices to pay the current clock price to all remaining active projects.

For the case in which costs are drawn from different distributions and/or project values differ, we restrict attention to the two-project case to retain tractability. In applications, the designer may prefer some projects over others and might have different information over cost distributions. In standard procurement settings, the quantity of units to be procured is not endogenously determined by a budget-constrained mechanism as in our model, but it is exogenously fixed to be some quantity k. It is well known that in k-unit procurement auctions the k projects with the greatest nonnegative virtual surpluses are implemented, e.g., Luton and McAfee (1986). In the asymmetric case, the ranking implied by costs and the ranking implied by virtual surpluses do not necessarily coincide. Broadly speaking, the designer discriminates against stochastically stronger projects, and prefers projects with higher values. The asymmetry requires that each project faces an individual clock and prices decrease asynchronously. In settings with exogenously given quantity restrictions, the clocks' speed can be optimally adjusted such that the virtual surplus of marginal projects is kept equal at all times, see Caillaud and Robert (2005, Proposition 1).

Interestingly, the optimality of such an allocation rule does not simply translate into the asymmetric case of our environment. In contrast, projects are not always greenlighted in order of their virtual surpluses. Therefore we cannot adopt the approach of Caillaud and Robert (2005). Instead, the descending-clock implementation of the optimal allocation includes individual clocks stopping at certain times. Here, a "quantity-quality tradeoff" kicks in: We show that the optimal allocation generically features instances in which out of two rival projects the project with lower virtual surplus is chosen. The reasoning behind this result is that the number of procured units is endogenous. In the asymmetric case, always greenlighting in order of virtual surplus reduces the expected number of greenlighted projects compared to the optimal mechanism. Strategyproofness creates a tradeoff between quantity (have a higher probability to implement more projects) and quality (guarantee to implement the superior project) of the procured projects. This discrimination of the stronger project is employed on top of the discrimination due to stochastic domination through the virtual costs.

Clock auctions are generally easy to understand and hard to manipulate. Furthermore, they are less information hungry than, for example, sealed bid auctions. In descending-clock auctions, the designer only learns the private information of those projects that are not greenlighted. In fact, Milgrom and Segal (2015) show that clock auctions are the only strategyproof mechanisms that preserve winners' unconditional privacy: Winners only need to reveal the minimum of their private information that is necessary to prove that they should be winning. These features of clock auctions make them attractive for applications in which there is limited trust between the involved parties.

To the best of our knowledge, this paper is the first that considers purely ex-post constrained optimal procurement design. Such a restrictive setting can be seen as a "worst-case scenario" for the designer, suiting many economic applications. In our leading example of the development fund, an ex-post budget constraint appears natural as budgets are usually fixed. The nonprofit nature of the projects might prohibit acquiring additional money on the financial market. Information rents are necessary, because a project might want to spend money on extra equipment that is convenient for the project's staff but has no value for the designer. In practice, such incentive problems are often resolved using dominant-strategy implementable mechanisms. In strategyproof mechanisms, agents have no incentive to invest in espionage activities or to hire consultants to avoid misspecification of beliefs. Mainly, dominant strategies are desirable as they are easy to explain and not prone to manipulation. For similar reasons, we restrict attention to deterministic mechanisms. Deterministic mechanisms obviate the need for a credible randomization device and are therefore more easily applicable in practice. As our agents care about their ex-post payoff, each constraint in a stochastic mechanism would have to hold for all outcomes of the mechanism's randomization anyway. Finally, ex-post participation constraints are necessary because projects simply cannot be conducted with insufficient funds, and the designer wants to avoid costly renegotiations when the projects default.

1.1 Literature

Even though the knapsack problem has a wide range of economic applications, there are relatively few publications in economics on this issue. Most prominently, Maskin (2002), in his Nancy L. Schwartz memorial lecture, addressed the related problem of the UK government that put aside a fixed fund to encourage firms to reduce their pollution. The government faces n firms that have private marginal cost of abatement θ_i and can commit to reduce x_i units of pollution. To reduce pollution as much as possible, the government pays expected compensation transfers t_i to the firms, who report costs and proposed abatement to maximize $t_i - \theta_i x_i$. For some distributions, Maskin (2002) proposes a mechanism that satisfies an ex-post participation constraint, an ex-post incentive compatibility constraint, and the condition that the budget is not exceeded in expectation. In response to Maskin (2002), Chung and Ely (2002b) look at a more general class of mechanism design problems with budget constraints and translate them into a setting à la Baron and Myerson (1982). Their approach nests Maskin (2002) and also Ensthaler and Giebe (2014a) as special cases. However, Ensthaler and Giebe (2014a) more explicitly derive a constructive solution. In contrast to us, they all consider a soft budget constraint that only requires the sum of expected transfers to be less than the budget. By incorporating the budget constraint into a Lagrangian function and ignoring the monotonicity (incentive) constraint, they find a mechanism that, under the standard regularity condition, indeed is incentive compatible.

In addition, Ensthaler and Giebe (2014a) use AGV-budget-balancing (such as Börgers and Norman, 2009) to obtain a mechanism which is ex-post budget-feasible. However, transformation into a mechanism with an ex-post balanced budget in such a way comes at the cost of sacrificing ex-post individual rationality. Many applications do not allow this constraint to be weakened. For instance, subsidy applicants usually cannot be forced to conduct their proposal when receiving only a small or possibly no subsidy. Alternatively, limited liability justifies insisting on ex-post individual rationality. Because we want both constraints to hold ex-post, we cannot build on their techniques and, thus, we approach the problem by characterizing the optimal allocation rule.

To the best of our knowledge, no paper exists that jointly considers optimal mechanism design under ex-post budget balance and ex-post individual rationality in a procurement setting. Ensthaler and Giebe (2014b) propose a belief-free clock mechanism that coincides with our optimal mechanism in the symmetric case for many parameterizations[4] but differs in the asymmetric case by holding the cost-benefit-ratio equal among projects. However, it has to be stressed that our mechanism designer knows the priors and projects' values, and exploits this knowledge, i.e., our mechanism is not detail-free. By simulating different settings, they conclude that this mechanism outperforms a mechanism used in practice. In contrast to their setting, the mechanism designer in our model values residual money. In Section 4, we discuss lesser weights on residual money and find that our main results qualitatively translates to the case in which residual money is neglected.

Because of the appeal of dominant-strategy incentive-compatible (DIC) mechanisms compared to Bayesian incentive-compatible (BIC) mechanisms, many researchers have produced valuable BIC-DIC equivalence results. These results characterize environments in which restricting attention to the more robust incentive criterion comes without loss. Our setup is not contained in these environments. For any BIC mechanism, Mookherjee and Reichelstein (1992) show that one can construct a DIC mechanism implementing the same ex-post allocation rule, whenever this allocation rule is monotone in each coordinate. However, the ex-post transfers of the constructed DIC mechanism are not guaranteed to satisfy ex-post budget balance. More recently, Gershkov, Goeree, Kushnir, Moldovanu, and Shi (2013) employ a definition of equivalence in terms of interim expected utilities introduced by Manelli and Vincent (2010). For any BIC mechanism, including the optimal one, they construct a DIC mechanism that yields the same interim expected utilities. Here, the ex-post allocation as well as the ex-post transfers might differ between the two. Therefore a DIC mechanism equivalent to a feasible BIC mechanism might violate the ex-post constraints in our setting.

Budget-constrained procurement setups have received much attention in the computer science literature. Instead of specifying the optimal mechanism, the authors in this literature typically aim to construct allocation algorithms that give good approximation guarantees. In other words, they try to maximize the minimal payoff an algorithm can guarantee compared to the full-information knapsack solver's payoff. Apart from the seminal paper by Singer (2010), the works of Dobzinski, Papadimitriou, and Singer (2011) and Chen, Gravin, and Lu (2011) are notable examples of this approach. Anari, Goel, and Nikzad (2014) present a stochastic algorithm and show that it gives the best possible approximation

[4]For all parameter constellations such that virtual surplus is always nonnegative.

guarantee in the many-projects limit in which any individual project's costs are small compared to the budget. While the above papers examine the belief-free case, Bei, Chen, Gravin, and Lu (2012) propose an algorithm for setups in which the designer knows how the private information is distributed.

Other auction theoretic papers featuring "knapsack auctions" deal with a slightly different problem compared to us. Aggarwal and Hartline (2006) consider a setting in which each agent is characterized by his object of commonly known size and a privately known valuation for having his object placed in the auctioneer's knapsack with commonly known capacity. They are looking for the truthful auction that best approximates the optimal full-information monotone pricing rule which maximizes the auctioneer's profit. Mu'Alem and Nisan (2008) cover the case of an auctioneer maximizing social welfare instead. Dütting, Gkatzelis, and Roughgarden (2014) study the performance of DA auctions for knapsack auctions, i.e., they show DA auctions fail to achieve a constant factor approximation of the optimal social welfare in knapsack auctions. Dizdar, Gershkov, and Moldovanu (2011) investigate a similar knapsack problem of a profit maximizing auctioneer in a dynamic setting: Agents sequentially arrive over time and are either included in the knapsack immediately or lost forever. Thereby they avoid combinatorial issues, which gives rise to a threshold property of the optimal mechanism. In such knapsack auctions, the mechanism designer maximizes the sum of transfers, and the value only enters the individual projects' payoff while the capacity constraint is imposed on the weight assigned to agents. In our framework, the value is collected by the auctioneer and the capacity constraint is imposed on the sum of transfers. Because of the latter, knapsack auctions and our knapsack procurement auctions are not dual problems.

There seems to be no reasonable analogy for our setting to another setting in which the mechanism designer is a similarly constrained seller and the agents are buyers. The literature on group-strategyproof cost-sharing mechanisms, initiated by Moulin (1999), considers the dual of a "surplus-sharing" problem. The crucial difference between this problem and our "budget-sharing" problem is that the agents themselves produce the output to be distributed, while in our case the budget to be distributed is fixed and unrelated to the surplus created by the agents, which is collected by the mechanism designer. Budget-constrained buyers in auctions have been discussed in the literature, e.g., by Che and Gale (1998) or Pai and Vohra (2014). However, these authors study budget-constrained agents whereas in our setting the designer is budget-constrained.

In the following section, we introduce the model. We start the analysis in Section 3 that is divided into a preliminary analysis for the general case, a full characterization of the general two-project case and the symmetric case, and finally a discussion of the general asymmetric case. Next, we discuss extensions and possible modifications to the model in Section 4. Finally, we conclude in Section 5.

2 Model

We consider a set of n projects $I = \{1, \ldots, n\}$ and one mechanism designer. The designer gains utility v_i if and only if project $i \in I$ is conducted. Each project can be conducted exactly once and its value is independent of the allocation.[5] We consider projects to be utility maximizing agents. If project i is executed, it incurs cost $c_i \in C_i := [\underline{c}_i, \overline{c}_i]$. Let $C := \times_{i \in I} C_i$ and $C_{-i} := \times_{j \in I \setminus \{i\}} C_j$. Let the realization of a cost vector be denoted by $\mathbf{c} = (c_i, \mathbf{c}_{-i}) \in C$. The costs are the projects' private information and are independently drawn from a distribution F_i. We assume F_i to be continuously differentiable with a strictly positive density f_i on the support. The value of the project v_i and the distribution F_i are common knowledge. If $F_i = F$ and $v_i = v$ for all $i \in I$, we refer to this environment as the symmetric case.

To compensate project i for its cost, the designer pays transfer t_i. We employ a revelation-principle argument and without loss of generality only consider direct mechanisms.[6] A direct mechanism is characterized by $\langle q_i, t_i \rangle$. It maps a vector of cost reports $\mathbf{c} \in C$ into binary provisions decision and transfers. We denote an allocation rule by $\gamma : C \to \mathcal{P}(I)$. It maps a cost vector into the set of "greenlighted" projects, an element of the power set of I. Correspondingly, we call $I \setminus \gamma(\mathbf{c})$ the set of "redlighted" projects, the projects that are not implemented.

We restrict attention to deterministic mechanisms. This restriction implies that

[5]That is, there are no exogenously given complementarities in a sense that one project's value increases when it is conducted together with another one or that the implementation of one project renders the other one worthless.

[6]In general, the revelation principle does not hold when restricting attention to deterministic mechanisms: Deterministic direct mechanisms are unable to replicate mixed-strategy equilibria in deterministic indirect mechanisms, as noted by, e.g., Strausz (2003). However, in our setting we do not lose generality. A mixed-strategy equilibrium consists of a distribution over pure-strategy profiles. Because the mechanism is implementable in dominant strategies any of these pure-strategy profiles also constitutes a pure-strategy equilibrium, in particular the pure-strategy equilibrium associated with the designer's most preferred outcome. Similarly, because the mechanism is ex-post constrained, this outcome is feasible. Therefore, while there are allocations that (in the class of deterministic mechanisms) can only be implemented by indirect mechanisms, the designer's most preferred feasible allocation can truthfully be implemented in a direct mechanism.

once all cost reports are collected, we know with certainty which project is selected by the mechanism. In other words, the decision of implementation q_i is binary,

$$q_i(\mathbf{c}) = \mathbb{I}(i \in \gamma(\mathbf{c})),$$

where $\mathbb{I}$ denotes an indicator function that equal one if the corresponding condi-
2 tion is true and zero otherwise.

Given a mechanism and a cost realization $\mathbf{c}$, project i's utility from reporting cost c_i' is given by its transfer minus the cost it bears,

$$u_i(c_i', \mathbf{c}) = t_i(c_i', \mathbf{c}_{-i}) - q_i(c_i', \mathbf{c}_{-i})c_i.$$

The designer derives value v_i from each greenlighted project i while having to pay the sum of transfers. Therefore she wants to maximize the aggregate value of greenlighted projects net of transfers paid. Her (ex-post) utility function u_D implies that the designer values residual money,

$$u_D(\mathbf{c}) = \sum_i \big(q_i(\mathbf{c})v_i - t_i(\mathbf{c})\big). \tag{1}$$

We impose an ex-post participation constraint. That is, if i is greenlighted the transfer must be at least as high as its cost,

$$t_i(c_i, \mathbf{c}_{-i}) - q_i(c_i, \mathbf{c}_{-i})c_i \geq 0 \quad \forall i \in I, (c_i, \mathbf{c}_{-i}) \in C. \tag{PC}$$

In addition, the designer has a budget constraint which is "hard" in the sense that she cannot spend more than her budget B for any realization of the cost vector. That is, the designer can never exceed her budget,

$$\sum_i t_i(\mathbf{c}) \leq B \quad \forall \mathbf{c} \in C. \tag{BC}$$

Finally, incentive compatibility has to hold ex-post. Alternatively, we can say that the mechanism has to be implementable in (weakly) dominant strategies[7] or that the mechanism must be strategyproof. Therefore for every realization of the cost vector, project i's truthful report must yield at least as much utility as any possible deviation,

$$\begin{aligned} t_i(c_i, \mathbf{c}_{-i}) - q_i(c_i, \mathbf{c}_{-i})c_i &\geq t_i(\widetilde{c}_i, \mathbf{c}_{-i}) - q_i(\widetilde{c}_i, \mathbf{c}_{-i})c_i \\ &\forall i \in I, \mathbf{c}_{-i} \in C_{-i} \text{ and } c_i, \widetilde{c}_i \in C_i. \end{aligned} \tag{IC}$$

[7] In our private value environment, these two concepts are equivalent in a direct revelation mechanism. In general, however, ex-post incentive compatibility is essentially a generalization of dominant-strategy implementability to interdependent value environments. See Chung and Ely (2002a).

3 Analysis

We search for the direct mechanism that maximizes the expected utility of the designer and refer to this mechanism as the optimal mechanism. One may think that a natural approach to this problem would be to express the ex-post transfer $t_i(c_i, \mathbf{c}_{-i})$ as a function of the ex-post allocation decision $q_i(c_i, \mathbf{c}_{-i})$, taking $\mathbf{c}_{-i}$ as given, and applying the envelope theorem. In that case, it would be possible to restrict attention to the allocation in order to solve for the optimal mechanism. However, this approach does not reduce the complexity of the problem. The reason is that the ex-post transfers and allocation for one cost vector restrict transfers and allocation for other cost vectors through the budget constraint in a manner much more involved than standard monotonicity. In particular, the budget constraint with the ex-post transfer expressed as a function of the ex-post allocation may be ill-behaved. Therefore we cannot straightforwardly arrive at sufficient conditions using convex optimization.[8]

Instead, we aim at deriving a set of properties that every mechanism must inherit to be optimal. We start by rewriting the general problem. For the case $n = 2$, we establish such properties by showing that the expected payoff yielded by any feasible mechanism not having one of the properties can be increased by adopting the properties. By virtue of these properties, the optimal allocation can be implemented by a myopic clock auction as defined by Milgrom and Segal (2015). These properties extend to the symmetric case. However, we provide asymmetric examples with $n > 2$ that violate the properties in the optimal mechanism, but are still implementable with a clock auction.

3.1 General preliminary analysis

Our first step is to show that strategyproofness implies that the optimal mechanism has to be a cutoff mechanism.

Lemma 1. *The optimal mechanism can be represented by cutoff functions $z_i : C_{-i} \to C_i$ such that project i is greenlighted if and only if it reports a cost weakly less than its cutoff,*

$$q_i(c_i, \mathbf{c}_{-i}) = \mathbb{I}(c_i \leq z_i(\mathbf{c}_{-i})).$$

The transfer to project i equals its cutoff if it is greenlighted and zero otherwise,

$$t_i(c_i, \mathbf{c}_{-i}) = q_i(c_i, \mathbf{c}_{-i}) z_i(\mathbf{c}_{-i}).$$

[8]Requiring either the budget or the participation constraint to hold only in expectation would enable us to use the techniques employed by Ensthaler and Giebe (2014a).

Proof. For any two cost reports $c_i, c_i' \in C_i$ of project i and for some $\mathbf{c}_{-i} \in C_{-i}$, (IC) implies that if the allocation of i is the same, $q_i(c_i, \mathbf{c}_{-i}) = q_i(c_i', \mathbf{c}_{-i})$, also the transfer has to be the same, $t_i(c_i, \mathbf{c}_{-i}) = t_i(c_i', \mathbf{c}_{-i})$. Otherwise, project i could, as one of the cost types, profitably deviate to the report yielding the higher transfer.

Conditional on i's allocation status and given any cost reports $\mathbf{c}_{-i}$, the transfer is fixed and does not vary with i's cost report. Hence, given $\mathbf{c}_{-i}$, there can only be two different transfers t_i for project i, one for each allocation status, $t_i^{q_i=1}(\mathbf{c}_{-i})$ and $t_i^{q_i=0}(\mathbf{c}_{-i})$.

Define $z_i(\mathbf{c}_{-i}) := t_i^{q_i=1}(\mathbf{c}_{-i}) - t_i^{q_i=0}(\mathbf{c}_{-i})$. Then, (IC) implies

$$q_i(c_i, \mathbf{c}_{-i}) = \begin{cases} 1 & \text{if } c_i \leq z_i(\mathbf{c}_{-i}) \\ 0 & \text{if } c_i > z_i(\mathbf{c}_{-i}) \end{cases}.$$

Suppose to the contrary that for some realization $\widehat{c}_i < z_i(\mathbf{c}_{-i})$ and some other $\widetilde{c}_i < z_i(\mathbf{c}_{-i})$, $q_i(\widehat{c}_i, \mathbf{c}_{-i}) = 0$ and $q_i(\widetilde{c}_i, \mathbf{c}_{-i}) = 1$. Then, type $\widehat{c}_i$ can profitably deviate to reporting $\widetilde{c}_i$ to ensure the green light which yields a utility increase of $z_i(\mathbf{c}_{-i}) - \widehat{c}_i$. An analogous argument applies for $\widehat{c}_i > z_i(\mathbf{c}_{-i}) > 0$.[9]

The last step is to show that $t_i^{q_i=0}(\mathbf{c}_{-i}) = 0$. This result follows from the mechanism being optimal, i.e., maximizing expected utility of the designer. □

As a direct consequence of dominant-strategy implementability, Lemma 1 shows that allocation and transfers are characterized by cutoffs. Project i is greenlighted whenever it reports a cost that lies weakly below the cutoff. Crucially, these cutoffs are functions of the other cost reports $\mathbf{c}_{-i}$. However, the optimal cutoffs remain to be determined. The maximization problem of the designer is

[9] When $c_i = z_i(\mathbf{c}_{-i})$, (IC) permits both $q_i(c_i, \mathbf{c}_{-i}) = 0$ and $q_i(c_i, \mathbf{c}_{-i}) = 1$. By convention, we assume $q_i(c_i, \mathbf{c}_{-i}) = 1$ in this case. However, writing a mechanism this way precludes the specification of tie-breakers, which might be necessary to conserve budget balance. For example, in a two-project example we would write down the mechanism "greenlight the cheaper project" as $z_1(c_2) = c_2$ and $z_2(c_1) = c_1$. If $c_1 = c_2$ a tie-breaker is needed to select a project. As this is a zero-probability event, the choice of the tie-breaker does not impact the designer's payoff. Similarly, as projects are indifferent, their ex-post utility is unaffected. Therefore we refrain from specifying a tie-breaker and proceed with our analysis as if both projects are greenlighted in these cases.

given by

$$\begin{gathered}\max_{\{z_i\}_{i\in I}} \mathbb{E}_{\mathbf{c}}\left[\sum_i q_i(\mathbf{c})v_i - t_i(\mathbf{c})\right] \\ \text{s.t. (BC)}, \\ q_i(\mathbf{c}) = \mathbb{I}(c_i \leq z_i(\mathbf{c}_{-i})) \quad \forall \mathbf{c} \in C, \\ t_i(\mathbf{c}) = \mathbb{I}(c_i \leq z_i(\mathbf{c}_{-i}))z_i(\mathbf{c}_{-i}) \quad \forall \mathbf{c} \in C.\end{gathered} \tag{2}$$

Incentive compatibility and participation constraints hold by construction, as q_i and t_i are determined by cutoff functions. Even the particularly crazy candidate in Figure 1 (introduced later) is incentive compatible and individually rational.

The next step towards solving this problem involves applying standard methods introduced by Myerson (1981). Let the conditional expected probability of being greenlighted and the conditional expected transfer be

$$\begin{aligned} Q_i(c_i) &= \mathbb{E}_{\mathbf{c}}[q_i(c_i, \mathbf{c}_{-i})|c_i] \\ \text{and } T_i(c_i) &= \mathbb{E}_{\mathbf{c}}[t_i(c_i, \mathbf{c}_{-i})|c_i]. \end{aligned}$$

The interim incentive compatibility required by Myerson (1981) is weaker than our condition (IC). Consequently, the expected transfer is determined by the allocation, $T_i(c_i) = Q_i(c_i)c_i + \int_{c_i}^{\bar{c}_i} Q_i(x)dx$. The usual monotonicity condition is trivially fulfilled as we are dealing with cutoff mechanisms. This reformulation in turn allows us to rewrite the objective function as a function of the allocation. Substituting into problem (2) and integrating by parts yields the following maximization problem,

$$\begin{gathered}\max_{\{z_i\}_{i\in I}} \mathbb{E}_{\mathbf{c}}\left[\sum_i \mathbb{I}(c_i \leq z_i(\mathbf{c}_{-i}))\left(v_i - c_i - \frac{F_i(c_i)}{f_i(c_i)}\right)\right] \\ \text{s.t.} \\ \sum_{i\in I} \mathbb{I}(c_i \leq z_i(\mathbf{c}_{-i}))z_i(\mathbf{c}_{-i}) \leq B \quad \forall \mathbf{c} \in C.\end{gathered} \tag{3}$$

We call $\varphi_i(c_i) := c_i + \frac{F_i(c_i)}{f_i(c_i)}$ the *virtual cost* of project i and $\psi_i(c_i) := v_i - \varphi_i(c_i)$ the *virtual surplus*. Here, φ and ψ are the procurement analogues to standard auction terminology. We can directly see from problem (3) that the optimal mechanism maximizes the expected sum of greenlighted virtual surpluses.

Note that constrained optimization by Lagrangian is not straightforward here because of the nondifferentiability of the indicator function. Instead, in the following we derive useful properties of the optimal cutoffs that can be exploited to characterize the optimal mechanism. A cutoff mechanism is by construction monotonic in the following sense:

Definition 1. *An allocation rule γ is monotonic in costs if $i \in \gamma(c_i, \mathbf{c}_{-i})$ and $c_i' < c_i$ imply $i \in \gamma(c_i', \mathbf{c}_{-i})$ for all $\mathbf{c}_{-i} \in C_{-i}$.*

In words, if a project gets greenlighted for some cost vector, it also gets greenlighted when, all else equal, its cost is lower. To proceed, we restrict the class of distributions from which costs can be drawn.

Assumption 1 (Log-concavity)**.** *For all i, the cumulative distribution function F_i is log-concave.*

This assumption is standard in information economics. It is equivalent to the reverse hazard rate function f/F being a weakly decreasing function or the ratio F/f being weakly increasing. Hence, the standard regularity condition is implied: φ_i is strictly increasing and ψ_i is strictly decreasing. A decreasing reverse hazard rate is the procurement analogue to the assumption of increasing hazard rate functions with a selling auctioneer.

Regularity ensures that a lower cost c_i translates to a higher virtual surplus $\psi_i(c_i)$. Hence, we can define the following cutoff cost type

$$z_i^{**} := \begin{cases} \psi_i^{-1}(0) & \text{if } \psi_i^{-1}(0) \in C_i \\ \bar{c}_i & \text{otherwise} \end{cases}, \tag{4}$$

where regularity implies the invertibility of ψ_i and thus allows for the above definition of z_i^{**}. In the symmetric case, $z_i^{**} = z^{**}$ for all $i \in I$. Let $\boldsymbol{\zeta}^{**}$ be the n-dimensional vector with z_i^{**} as i-th element for all $i \in I$.

Definition 2. *An allocation rule γ is $\boldsymbol{\zeta}^{**}$-exclusive if, for all $i \in I$, $c_i > z_i^{**}$ implies that $i \notin \gamma(c_i, \mathbf{c}_{-i})$ for all $\mathbf{c}_{-i} \in C_{-i}$.*

A cutoff mechanism is $\boldsymbol{\zeta}^{**}$-exclusive if and only if $z_i(\mathbf{c}_{-i}) \leq z_i^{**}$ for all $\mathbf{c}_{-i} \in C_{-i}$ and for all $i \in I$. If the budget sufficed, a designer would want to greenlight all projects with nonnegative virtual surplus. Crucially, the arguments leading to this statement also imply that it is never optimal to greenlight a project with negative virtual surplus.

Lemma 2. *The optimal mechanism is $\boldsymbol{\zeta}^{**}$-exclusive. In the trivial case, $\sum z_i^{**} \leq B$, the optimal cutoffs are independent of the cost reports,*

$$z_i(\mathbf{c}_{-i}) = z_i^{**} \quad \forall \mathbf{c}_{-i} \in C_{-i} \text{ and } \forall i \in I.$$

The proof of this lemma is standard and hence omitted. It immediately follows from the rewritten objective function (3): Greenlighting a project with negative virtual surplus decreases the designer's payoff and uses part of the budget. Guaranteeing the green light for high-cost types comes at the cost of having to pay higher information rents to all cost types. For the same reason, also a budget-unconstrained designer would implement a $\boldsymbol{\zeta}^{**}$-exclusive mechanism, even when the surplus $v_i - \bar{c}_i$ is positive for all projects.

To continue our analysis, we focus on the more tractable two-project case in the next subsection. The first aim is to provide more structure on the cutoff functions that determine the optimal allocation. This enables us to fully characterize the optimal allocation and how to implement it.

3.2 $n = 2$

Reducing the set of mechanisms that are candidates for optimality implies a strong property in the two-project case: Project substitutability means that, if a project gets greenlighted for some cost vector $\mathbf{c}$, it is also greenlighted when, all else equal, another project's cost is increased. This property relates to the cross-monotonicity defined in the cost-sharing problem of Moulin and Shenker (2001): An agent's cost share cannot increase when the allocation set expands.

Definition 3. *An allocation rule γ has substitutes if $i \in \gamma(\mathbf{c})$ and $c'_j > c_j$ for some $j \neq i$ implies $i \in \gamma(c'_j, \mathbf{c}_{-j})$. Otherwise, the allocation rule has complements.*

Having in mind a setting with an exogenously determined amount of projects to be procured and without a budget constraint, this property is clearly optimal, because if i is among the projects with the highest virtual surpluses for some cost vector, it is also among them when the cost of some other project j is increased, i.e., when j's virtual surplus is decreased. However, with the budget constraint, this property does not hold in a full-information setting.[10] A cutoff mechanism has substitutes if all functions z_i are weakly increasing in each argument.

The optimality of this property will be proved jointly with Lemma 3 and Lemma 4. The intuition is straightforward: The cost realizations of all projects are independent and therefore project i's cost report only influences the allocation of project $j \neq i$ via the budget constraint. Project i's cost report fixes the cutoff of project j and thus determines the residual budget available to itself in

[10] For example, there are two projects, $v_1 > v_2$. Under full information, both projects get implemented for a cost vector $(c_1, c_2) = (B - z, z)$. Then, increasing c_1 would kick project 2 out of the allocation. In contrast, in our asymmetric-information setting where c_2 pins down a cutoff $z_1(c_2)$ for project 1, project 1 instead loses the green light status, when its cost increases while c_2 remains constant.

case project j is implemented. If project i exceeds its cutoff, this frees budget to be distributed to project j. Consequently, j's transfer should either remain constant or increase. For $n > 2$, a cost report does not simultaneously pin down all other cutoffs and the remaining budget. In asymmetric cases, it is possible that projects endogenously become complements, see Subsection 3.4.

Project substitutability is related to the next property, non-bossiness. However, the two properties are not equivalent since an allocation rule can have non-bossy complements as seen in Example 2. Milgrom and Segal (2015) provide an example with substitutes and a bossy winner.

Definition 4. *An allocation rule γ is non-bossy if $\gamma(c_i', \mathbf{c}_{-i}) \cap \{i\} = \gamma(\mathbf{c}) \cap \{i\}$ implies $\gamma(c_i', \mathbf{c}_{-i}) = \gamma(\mathbf{c})$.*

An allocation rule γ has non-bossy winners if for any $i \in I$, $\mathbf{c} \in C$, and $c_i' \in C_i$, $i \in \gamma(c_i', \mathbf{c}_{-i}) \cap \gamma(\mathbf{c})$ implies $\gamma(c_i', \mathbf{c}_{-i}) = \gamma(\mathbf{c})$. Otherwise, winners can be bossy.

An allocation rule γ has non-bossy losers if for any $i \in I$, $\mathbf{c} \in C$, and $c_i' \in C_i$, $i \notin \gamma(c_i', \mathbf{c}_{-i}) \cup \gamma(\mathbf{c})$ implies $\gamma(c_i', \mathbf{c}_{-i}) = \gamma(\mathbf{c})$. Otherwise, losers can be bossy.

An allocation has non-bossy substitutes if it has substitutes and is non-bossy. An allocation has non-bossy complements if it has complements and is non-bossy.

In words, a non-bossy winner (loser) cannot affect the allocation without changing its own green-light (red-light) status. In Example 3, we illustrate that an optimal allocation rule can have bossy losers when there are at least three projects. The following lemma states that, given both projects are greenlighted for two different cost vectors, the transfers for both cost vectors have to be the same. That is, when both projects are greenlighted, their transfer is constant. Intuitively, optimal cutoffs cannot depend on greenlighted projects' cost, because for these projects the cutoff coincides with the transfer. If the budget constraint is binding, a greenlighted project would be able to influence its own cutoff, i.e., the budget minus the transfer to the other (greenlighted) project. This contradicts the notion of a cutoff mechanism.

Lemma 3. *Suppose the nontrivial case with $n = 2$. If γ is optimal and $\gamma(\widetilde{c}_1, \widetilde{c}_2) = \gamma(\widehat{c}_1, \widehat{c}_2) = \{1, 2\}$, the transfers to both projects are constant. That is,*

$$\begin{aligned} t_1(\widetilde{c}_1, \widetilde{c}_2) &= z_1(\widetilde{c}_2) = z_1(\widehat{c}_2) &= z, \\ t_2(\widetilde{c}_1, \widetilde{c}_2) &= z_2(\widetilde{c}_1) = z_2(\widehat{c}_1) &= B - z. \end{aligned}$$

Proof. By Lemma 1, the optimal mechanism has to be a $\boldsymbol{\zeta}^{**}$-exclusive cutoff mechanism. Take any feasible candidate mechanism with any cutoff functions

$\{z_1, z_2\}$ violating the lemma and define

$$\begin{aligned} a_1 &= \max\{c_1 | \exists c_2 : c_2 \leq z_2(c_1), c_1 \leq z_1(c_2)\} \\ a_2 &= \max\{c_2 | \exists c_1 : c_1 \leq z_1(c_2), c_2 \leq z_2(c_1)\}, \end{aligned} \tag{5}$$

i.e., a_i is the highest cost of project i such that both projects are implemented. Since by assumption there exist cost vectors such that both projects are greenlighted, the sets over which we have defined a_1 and a_2 are non-empty. The maximum exists by left-continuity of any optimal function z_i.[11] For a graphical illustration of a feasible candidate mechanism violating the lemma and of how to improve such a mechanism, consult Figure 1.

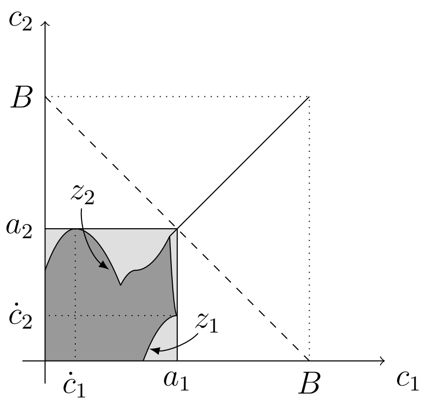

Figure 1: The depicted feasible mechanism greenlights both project for all cost combinations in the darker gray area. This mechanism cannot be optimal, since the alternative mechanism constructed is feasible as well, additionally greenlights a project in the lighter gray area and is otherwise equivalent.

Hence by definition of a_1, there (not necessarily uniquely) exists $\dot{c}_2$ such that $a_1 = z_1(\dot{c}_2)$. Similarly, there exists $\dot{c}_1$ such that $a_2 = z_2(\dot{c}_1)$. By definition, $(\dot{c}_1, \dot{c}_2) \leq (a_1, a_2)$ and at cost realization $(\dot{c}_1, \dot{c}_2)$ both projects are implemented. The budget feasibility of the candidate mechanism implies $a_1 + a_2 \leq B$ such that the following constructed alternative mechanism is feasible as well.

The initial candidate cannot be optimal, since it is outperformed by an alterna-

[11] We can replace any function z_i with a left-continuous function that is identical up to a set of points with Lebesgue-measure zero. Hence, if there exists an optimal function z_i that is not left-continuous, then there also exists a left-continuous version of the same function that yields the same payoff and hence is also optimal.

tive mechanism with cutoffs

$$\widehat{z}_i(c_j) = \begin{cases} a_i & \text{if } c_j \leq a_j \\ z_i(c_j) & \text{else.} \end{cases}$$

This alternative mechanism weakly outperforms the initial candidate state-by-state as it either implements the same allocation or a strictly better one by greenlighting an additional project. If $a_1 + a_2 \neq B$ and both $a_i < z_i^{**}$, then the alternative mechanism can be improved further by increasing cutoffs such that the budget constraint binds. □

In Lemma 5, a form of the previous lemma is generalized to $n > 2$ if γ has substitutes: Given two cost vectors implement the same allocation and only differ in the cost levels of greenlighted projects, the transfers to these greenlighted projects are identical for both cost vectors.

Lemma 4. *Suppose $n = 2$. The optimal mechanism has substitutes,*

$$z_i(\widetilde{c}_j) \geq z_i(\widehat{c}_j) \quad \text{for almost every } \widetilde{c}_j > \widehat{c}_j \tag{6}$$

and is non-bossy.

Proof. In Lemma 3, we have shown that an optimal z_i is constant for all $c_j \leq a_j$, as defined in (5). Consequently, the set of cost combinations such that both projects are implemented is

$$A := \{(c_1, c_2) : c_1 \leq a_1, c_2 \leq a_2\}$$

and $|\gamma(c_1, c_2)| \leq 1$ if $(c_1, c_2) \notin A$. By the cutoff nature of the optimal mechanism, we also know that $j \notin \gamma(c_1, c_2)$ if $c_i \leq a_i$ and $c_j > a_j$. By Lemma 2, $\gamma(c_1, c_2) = \emptyset$ if $(c_1, c_2) \in Z^{**}$ with

$$Z^{**} := \{(c_1, c_2) : c_i > \min\{z_i^{**}, B\} \text{ for both } i \in \{1, 2\}\}.$$

As a result, it is feasible and optimal to greenlight exactly one project for cost vectors $(c_1, c_2) \notin A \cup Z^{**}$. If it is optimal and feasible to implement project 1 given some cost vector $(z_1(\widehat{c}_2), \widehat{c}_2)$, then, by regularity, the mechanism designer also prefers to greenlight project 1 when the cost of project 2 is increased to $\widetilde{c}_2 > \widehat{c}_2$ with

$$\psi_1(z_1(\widehat{c}_2)) \geq \psi_2(\widehat{c}_2) > \psi_2(\widetilde{c}_2), \tag{7}$$

i.e., decreasing cutoffs are suboptimal, $z_1(\widetilde{c}_2) \geq z_1(\widehat{c}_2)$. Since cutoff functions are weakly increasing, a project 2 cannot kick project 1 out of the allocation by reporting a higher cost. If project 2 reports a lower cost, it either does not change the allocation (if $c_2 \leq a_2$) or project 1 is replaced by project 2, because the first inequality in (7) is flipped. As a result, winners and losers are non-bossy. □

On a first glance, Inequalities (7) seem to suggest that it is always optimal to greenlight the project with the higher virtual surplus when $|\gamma(c_1, c_2)| = 1$. However, always greenlighting the better project may not be feasible when the set A is determined optimally. Suppose that in optimum, $\psi_1(a_1) > \psi_2(a_2)$. Then, an allocation $\gamma(a_1 + \varepsilon, a_1) = 1$ for some $\varepsilon > 0$ would not be strategyproof: Given $c_2 = a_2$, project 1 of cost type a_1 would want to misreport costs to gain a transfer of at least $a_1 + \varepsilon$. Such a combination of (a_1, a_2) turns out to be a generic feature of the optimal mechanism.

Let us consider the nontrivial two-project case: The rewritten maximization problem of the designer (3) is given by

$$\begin{gathered}\max_{z_1(c_2),z_2(c_1)} \mathbb{E}\Big[\mathbb{I}(c_1 \le z_1(c_2)) \left(v_1 - c_1 - \tfrac{F_1(c_1)}{f_1(c_1)}\right) \\ +\mathbb{I}(c_2 \le z_2(c_1)) \left(v_2 - c_2 - \tfrac{F_2(c_2)}{f_2(c_2)}\right)\Big] \\ \text{s.t.} \\ \mathbb{I}(c_1 \le z_1(c_2))z_1(c_2) + \mathbb{I}(c_2 \le z_2(c_1))z_2(c_1) \le B \quad \forall (c_1, c_2) \in C.\end{gathered} \tag{8}$$

By virtue of the optimal properties, the designer must greenlight project i once its cost is below the constant a_i. If both projects report greater costs, the designer is free to choose one of them. A glance at the objective function (8) reveals that in such a case it is desirable to greenlight the project with greater positive virtual surplus. This insight allows us to rewrite the objective function (8) as a function of constant z,

$$\begin{aligned}\max_z \pi(z) = &\int_0^z \psi_1(c_1) dF_1(c_1) + \int_0^{B-z} \psi_2(c_2) dF_2(c_2) \\ &+ \int_{\max\{\psi_2^{-1}(\psi_1(z)),B-z\}}^{\bar{c}_2} \int_z^{\min\{\psi_1^{-1}(\psi_2(c_2)),z_1^{**},B\}} \psi_1(x) dF_1(x) dF_2(c_2) \\ &+ \int_{\max\{\psi_1^{-1}(\psi_2(B-z)),z\}}^{\bar{c}_1} \int_{B-z}^{\min\{\psi_2^{-1}(\psi_1(c_1)),z_2^{**},B\}} \psi_2(x) dF_2(x) dF_1(c_1),\end{aligned} \tag{9}$$

i.e., the problem collapses to finding a single constant.

In the symmetric case, the ranking of virtual surpluses coincides with the reversed order of costs. Hence, it can easily be seen that in the symmetric case either both projects are implemented or the one with lower costs. This observation for the symmetric case extends to more than two projects, $n > 2$, see Subsection 3.3. A natural extension of this mechanism to the asymmetric case would involve adjusting the cutoffs so that they equalize virtual surplus. This modification ensures that, if a project has to be redlighted, the least attractive project in terms of virtual surplus is rejected. We call this allocation rule the candidate allocation.

The nongeneric condition for optimality of the candidate allocation is stated in (10). To implement the candidate allocation, the constant cutoffs at which both projects are greenlighted must be a pair $(a_1, a_2) = (z, B - z)$ such that $\psi_1(z) = \psi_2(B-z)$. Then, however, optimality is only obtained if $\frac{F_2(B-z)}{f_2(B-z)} = \frac{F_1(z)}{f_1(z)}$. The intuition behind this statement is straightforward. Selecting z in order to satisfy $\psi_1(z) = \psi_2(B - z)$ allows the designer to always greenlight the project with the higher virtual surplus, whenever it is not feasible to greenlight both projects. However, if $\frac{F_2(B-z)}{f_2(B-z)} \neq \frac{F_1(z)}{f_1(z)}$ the cutoffs z and $B - z$ do not maximize the probability to greenlight both projects. Consequently, the designer can adjust the cutoffs $\{z, B - z\}$ to trade off a higher probability of implementing the most favorable allocation ($\gamma(c_1, c_2) = \{1, 2\}$) against a positive probability of having to implement the less preferred of two possible singleton allocations ($\gamma(\mathbf{c}) = \{j\}$, when project j has lower virtual surplus).

Therefore two aspects of the designer's payoff maximization - getting projects with high virtual surplus and getting as many projects as possible - are only aligned if condition (10) is met. In the symmetric case, the condition holds by construction. However, in an asymmetric environment it is generically violated.

Proposition 1. *In the nontrivial asymmetric two-project case, i.e., $n = 2$ and $z_1^{**} + z_2^{**} > B$, in which values or cost distributions differ across projects, it is generically not optimal to always greenlight the project with the higher virtual surplus. That is, under the optimal allocation rule γ, there may exist cost vectors such that*

$$i \notin \gamma(c_i, c_j), \text{ and } j \in \gamma(c_i, c_j)$$

although

$$\psi_i(c_i) > \psi_j(c_j).$$

Proof. Given the max operators in (9), the derivative takes a different form depending on whether $\psi_1(z) \gtrless \psi_2(B - z)$. However, as π is continuously differentiable, it suffices to look at one of the two forms,

$$\begin{aligned}\left.\frac{\partial \pi}{\partial z}\right|_{z:\psi_1(z)\geq\psi_2(B-z)} &= \int_z^{\psi_1^{-1}(\psi_2(B-z))} \psi_1(x) dF_1(x) f_2(B-z) + \\ &+ \psi_1(z) f_1(z) F_2(B-z) \\ &- \psi_2(B-z) f_2(B-z) F_1(\psi_1^{-1}(\psi_2(B-z))).\end{aligned}$$

Now, consider z corresponding to the candidate allocation with $\psi_1(z) = \psi_2(B - z)$, which yields

$$\frac{\partial \pi}{\partial z} = 0 \Leftrightarrow \frac{F_2(B-z)}{f_2(B-z)} = \frac{F_1(z)}{f_1(z)}, \tag{10}$$

a nongeneric case. Consequently, it is generically not optimal to always allocate to the project with the higher virtual surplus. □

Proposition 1 is driven by a tradeoff between quantity and quality: Even though the designer prefers the project with higher virtual surplus conditional on implementing only a single project, she sometimes greenlights the project with lower virtual surplus out of two rival projects, as quantity is endogenous here. By endogenous quantity, we mean that the designer is only restricted by the feasibility constraints and is otherwise free to choose how many projects she want to procure. The simplest way to lay out the intuition behind Proposition 1 is by an example.

Example 1. *There are two projects,* $(n = 2)$ *with* $v_1 = 5, v_2 = 4.5$ *and* c_1 *and* c_2 *are uniformly distributed on support* $[0, 1]$. *The budget is given by* $B = 1$. *The optimal cutoff functions are given by:*

$$z_1(c_2) = \begin{cases} 0.53 & \text{if } c_2 \leq 0.47 \\ c_2 + 0.25 & \text{if } 0.47 < c_2 \leq 0.75 \\ 1 & \text{if } c_2 > 0.75 \end{cases}$$

$$z_2(c_1) = \begin{cases} 0.47 & \text{if } c_1 \leq 0.72 \\ c_1 - 0.25 & \text{if } c_1 > 0.72. \end{cases}$$

The corresponding optimal allocation is:

$$(q_1, q_2) = \begin{cases} (1,1) & \text{if } 0 \leq c_1 \leq 0.53 \text{ and } 0 \leq c_2 \leq 0.47 \\ (1,0) & \text{if } 0 \leq c_1 \leq 0.72 \text{ and } c_2 > 0.47 \\ (1,0) & \text{if } c_1 > 0.72 \text{ and } \psi_1 \geq \psi_2 \\ (0,1) & \text{if } 0.53 < c_1 \leq 0.72 \text{ and } c_2 \leq 0.47 \\ (0,1) & \text{if } c_1 > 0.72 \text{ and } \psi_1 < \psi_2. \end{cases}$$

The corresponding transfers are:

$$t_1(c_1,c_2) = \begin{cases} 0.53 & \text{if } c_2 \le 0.47 \text{ and } c_1 \le 0.53 \\ c_2 + 0.25 & \text{if } 0.47 < c_2 \le 0.75 \text{ and } c_1 \le c_2 + 0.25 \\ 1 & \text{if } c_2 > 0.75 \\ 0 & \text{otherwise} \end{cases}$$

$$t_2(c_1,c_2) = \begin{cases} 0.47 & \text{if } c_1 \le 0.72 \text{ and } c_2 \le 0.47 \\ c_1 - 0.25 & \text{if } c_1 > 0.72 \text{ and } c_2 < c_1 - 0.25 \\ 0 & \text{otherwise.} \end{cases}$$

Consider Example 1. The candidate allocation demands cutoffs such that $\widetilde{a_1} = 0.625$ and $\widetilde{a_2} = 0.375$ for allocating to both projects. At these cutoffs, the probability of greenlighting both projects is $0.625 \cdot 0.375 \approx 0.234$. This allocation is depicted in Panel 2a. In contrast, the maximal feasible probability to greenlight both projects is at equal cutoffs, $\widehat{a_1} = \widehat{a_2} = 0.5$. The corresponding area is the dotted square in the lower-left corner of Panel 2b. However, at these cutoffs it is not incentive compatible to guarantee the green light for the project with higher virtual surplus in every case. More specifically, it is not incentive compatible to allocate along the dotted[12] diagonal line, if at least one project exceeds $\widehat{a_i}$. Hence, strategyproofness introduces a tradeoff between maximizing the probability of greenlighting both projects and allocating to the preferred one if only one project is feasible. Consequently, the optimal (a_1^*, a_2^*) do not lie at $(0.625, 0.375)$ but rather at $(0.53, 0.47)$. Importantly, this optimal discrimination against the stronger project is pursued on top of the usual discrimination against stochastically stronger projects reflected in the virtual costs.

Given the optimal allocation in Example 1, there are some realizations of the cost vector for which the designer greenlights the project with lower virtual surplus. These realizations are represented by the shaded area in Panel 3a. Here, the constraints and the choice of (a_1, a_2) force the designer to greenlight project 2, even though project 1 has the higher virtual surplus.

The cost vectors for which the designer implements both projects are represented by the rectangular area in the lower-left corner of Panel 3a. Any point (a_1, a_2) on the dashed line representing the budget constraint satisfies $a_1 + a_2 = B$.

[12]Not to be confused with the dashed diagonal representing the budget constraint.

Moving this corner point southeast from $(0.5, 0.5)$ along the dashed budget line has two effects: shrinking the shaded area and shrinking the area of the lower-left rectangle. While it is desirable to shrink the shaded area, in which the designer must allocate to project 2 despite its lower virtual surplus, it is undesirable to shrink the size of the rectangle, which in this example represents the probability that both projects are conducted. Given that we have an interior solution in this example, at (a_1, a_2) these two effects balance each other out.

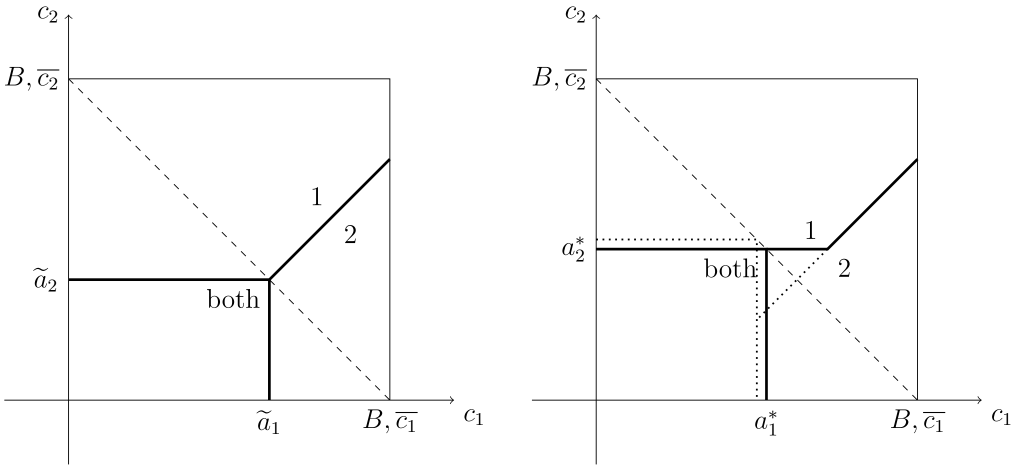

(a) Candidate allocation.

(b) Optimal allocation.

Figure 2: Candidate and optimal allocation for Example 1.

Graphically, the fact that there is no slack in the budget constraint if both projects are greenlighted implies that the area representing points at which both projects are executed touches the dashed line at least once. In fact, it can touch the budget line exactly once, as it is not possible to greenlight both projects when $c_1 > a_1$ or $c_2 > a_2$ without violating (BC) sometimes. This result means that the area where both projects are greenlighted is the rectangle with corners $(0, 0)$ and (a_1, a_2). Then, if $c_1 < a_1$ but $c_2 > a_2$, the nature of cutoffs prevents the designer from greenlighting project 2. Therefore project 1 must be greenlighted, as represented by the lightly shaded area in Panel 3b. A similar argument applies to the darkly shaded area. Thus, looking at Panel 3b, the choice of (a_1, a_2) determines the allocation for all cost realizations except those in the upper-right corner. Here, the designer is free to choose the allocation, as long as the line delineating whether project 1 or 2 gets greenlighted is (weakly) increasing or vertical. Not surprisingly, it is optimal to greenlight the project with the higher virtual surplus.

Having characterized the optimal allocation, we now turn to the issue of how to

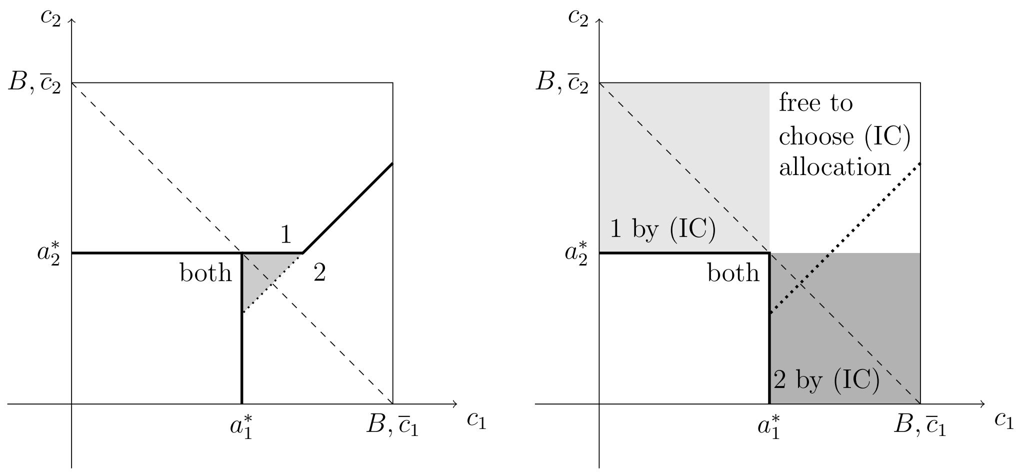

(a) Greenlighting the project with lower virtual surplus.

(b) (IC)-constraints on the allocation.

Figure 3: Greenlighting the project with lower virtual surplus and (IC)-constraints on the allocation (Example 1).

implement it. Taking stock, among all mechanisms satisfying (PC), (BC) and (IC), any mechanism that maximizes the designer's expected payoff (1) belongs to a certain class of mechanisms: We have shown that the optimal two-project mechanism is

Property 1 monotonic in costs,

Property 2 $\boldsymbol{\zeta}^{**}$-exclusive,

Property 3 non-bossy, and

Property 4 has substitutes.

Being able to restrict attention to mechanisms with these properties is highly useful, as these mechanisms are a much more tangible class than the substantially larger set of all permissible cutoff mechanisms. In addition, all mechanisms with these properties can be implemented with a DA auction as proposed by Milgrom and Segal (2015). To this end, we first restate their definition adapted to our setting.

Definition 5 (DA auction). *A deferred acceptance (DA) auction is an iterative algorithm defined by a collection of scoring functions*

$$s_i^A : C_i \times C_{I \setminus A} \to \mathbb{R}_+$$

that are weakly increasing in c_i for all $i \in A$ and for all $A \subset I$. Let $A_t \subset I$ denote the set of active bidders in iteration t and initially $A_1 = I$. The algorithm stops in some period T when all active projects have a score of zero, $s_i^{A_T} = 0$ for all $i \in A_T$. Then the set of greenlighted project is A_T. Otherwise, at each iteration t, the project with the highest score is removed. The payment p_i^t of project i at iteration t is either given by the highest possible cost that i could have had without being removed from the set of active bidders or by the last iteration's payment, depending on which payment is smaller,

$$p_i^t(\mathbf{c}) = \begin{cases} \sup\{c_i' : s_i^{A_t}(c_i', \mathbf{c}_{I\setminus A_t}) < s_j^{A_t}(c_j, \mathbf{c}_{I\setminus A_t})\} & \text{for } j \in A_t \setminus A_{t+1}, \\ \min\{\sup\{c_i' : s_i^{A_t}(c_i', \mathbf{c}_{I\setminus A_t}) \leq 0\}, p_i^{t-1}\} & \text{if } t = T. \end{cases}$$

The algorithm is initialized with $p_i^0 = \min\{\bar{c}_i, z_i^{**}, B\}$.[13]

The main appeal of DA auctions lies in their incentive guarantees. They are not only strategyproof, they are obviously strategyproof, as defined by Li (2015). Moreover, DA auctions are weakly group-strategyproof. That is, no coalition of projects can manipulate their reports such that it strictly increases the utility of all projects in the coalition: At least one member of the coalition receives a weakly worse payoff whenever other coalition members benefit. Because collusion in auctions is generally illegal, compensating the worse off coalition member is not contractible. In addition, the dominant-strategy equilibrium outcome in a DA auction can be interpreted as robust in the following sense: Consider the full-information game in which all cost reports are observed, projects can report any cost, the allocation is determined according to the DA auction's allocation rule, but projects receive their own report as payments. The dominant-strategy equilibrium outcome of the DA auction is the only outcome that survives iterated deletion of dominated strategies in this game.

Proposition 2. *(Milgrom and Segal (2015)) Any monotonic allocation rule with substitutes and non-bossy winners has a DA auction representation and can be implemented with a descending-clock auction.*

Milgrom and Segal (2015) prove this statement for finite type spaces in their Proposition 7. The informed reader may notice that their proposition is an if-and-only-if statement. However, the necessity of the substitutes condition hinges on the fact that they want the statement to hold for any subset of the type space. We discuss this necessity further in Subsection 3.4. By Proposition 2,

[13] Compared to Milgrom and Segal (2015), we slightly tweak the updating function of payments without changing the deferred acceptance nature of the algorithm and any of its properties.

the optimal allocation can be implemented with a descending-clock auction. In the following, we show how to accommodate the tradeoff between quantity and quality elaborated on in Proposition 1 in a modified clock auction.

Corollary 1. *Generically, in an optimal implementation with descending price clocks, the clocks not only run at individual speeds, occasionally some clocks also have to halt.*

Since the optimal allocation rule in the symmetric case is anonymous, it can be implemented with a single clock that suggests prices for all active projects. However, in asymmetric cases, each project must have an individual price clock, because heterogeneous virtual surplus functions require individual speeds. Interestingly, an implication of the quantity-quality tradeoff is that sometimes one clock has to halt. For Example 1, the clock prices, denoted by τ_i, are depicted in Figure 4 as a function of time. The duration of the auction can be divided into three segments. The auction starts with both clocks at $z_1^{**} = z_2^{**} = \overline{c}$. First, τ_2 decreases while τ_1 is held constant, which happens until both clock prices lead to the same virtual surplus, i.e., $\psi_2(\tau_2) = \psi_1(\overline{c}_2)$. Second, both τ_1 and τ_2 decrease simultaneously, but asynchronously keeping virtual surplus equal, $\psi_1(\tau_1) = \psi_2(\tau_2)$, until $\tau_2 = z_2(\underline{c}_1)$. Third, only τ_1 decreases until $\tau_1 = z_1(\underline{c}_2)$. If at this point both projects still remain in the auction, the auction stops and both are greenlighted. Otherwise, the inferior project 2 is greenlighted.

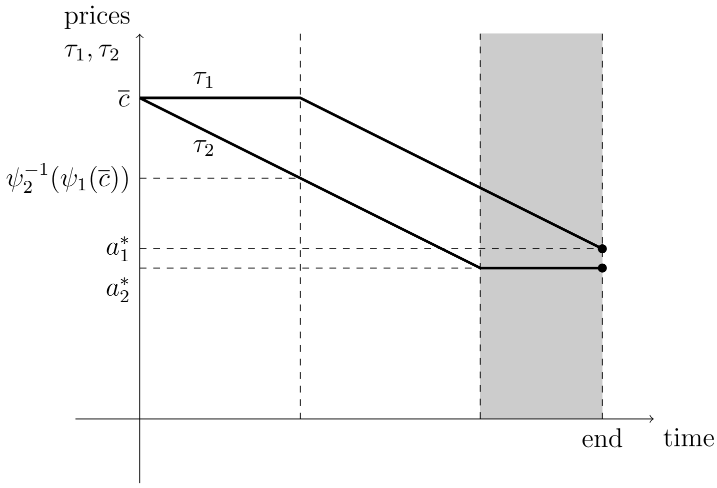

Figure 4: Optimal descending-clock auction in Example 1.

The cost vectors for which the designer greenlights project 2 despite its lower virtual surplus, represented by the shaded area in Panel 3a, are also represented graphically in Figure 4: If the auction ends in the third time segment (shaded area of Figure 4) before both projects can be greenlighted, project 1 must have exited because τ_1 dropped below c_1. Project 2 is greenlighted and receives transfer a_2^* even though project 1 has the higher virtual surplus. Therefore if cost vectors in the shaded area of Panel 3a realize, the optimal descending-clock auction ends in the third time segment.

We should emphasize again a novel feature of this descending-clock auction. The clocks of both projects are paused asynchronously over some time of the auction. One project's clock runs down while the other project's clock stops. Since we have examined a very simple example with only two projects, each project's clock is paused only once. If an allocation with more projects is implementable with price clocks, the projects' clocks may pause and resume several times.

Given the complexity of our problem, we do not find a simple and general ($n > 2$) full characterization of the optimal mechanism in the asymmetric case that we further elaborate on in Subsection 3.4. In our examples with two projects, the problem boils down to finding one point, (a_1, a_2), with respect to one crucial tradeoff. Naturally, the number of relevant tradeoffs increases with the number of projects. Therefore unfortunately, optimization with a larger set of projects quickly loses tractability.

3.3 $n > 2$. The symmetric case

In this section, we focus on symmetric projects, i.e., environments with $v_i = v$ and $F_i = F$ for every project $i \in I$. An implication of this assumption is that the order of costs coincides with the order of virtual surpluses and that $z_i^{**} = z^{**}$ for all $i \in I$. Hence, there is no such tradeoff as in asymmetric cases: The designer can maximize the probability to implement the best allocation (greenlight as many projects as the budget allows) without being forced to greenlight an inferior project by the incentive constraint.

Proposition 3. *Arrange the projects in ascending order of their reported costs, $c_1 \leq c_2 \leq \cdots \leq c_n \leq c_{n+1} := \bar{c}$, and define $z^k := \min\left\{\frac{B}{k}, z^{**}, c_{k+1}\right\}$. In the symmetric case, given any cost vector, the optimal number of implemented projects k^* is given by $k^* := \max\{k | c_k \leq z^k\}$ and all implemented projects receive identical transfer z^{k^*}.*

Proof. First of all, define $\overline{k} = \max\{k : k\underline{c} \leq B\}$ as the maximal possible number of procured projects. Even under full information it is never budget-feasible to implement more than $\overline{k}$ projects. Consequently, implementing the cheapest $\overline{k}$

projects is the designer's most favorable allocation. Let $c_{-i,k}$ be the k-th highest cost of i's competitors. By setting

$$z_i(\mathbf{c}_{-i}) = \min\{z^{**}, c_{-i,\overline{k}}\} \quad \text{if } c_{-i,\overline{k}} \leq \frac{B}{\overline{k}} \tag{11}$$

for all $i \in I$ the designer guarantees the first-best allocation for all vectors such that $c_k \leq \frac{B}{\overline{k}}$ for any $k > \overline{k}$: The cheapest $\overline{k}$ projects are greenlighted when they have nonnegative virtual surplus, otherwise all projects with nonnegative virtual surplus are greenlighted.

By setting

$$z_i(\mathbf{c}_{-i}) = \min\{z^{**}, \frac{B}{\overline{k}}\} \quad \text{if } c_{-i,\overline{k}-1} \leq \frac{B}{\overline{k}} < c_{-i,\overline{k}} \tag{12}$$

for all $i \in I$ the designer guarantees the first-best allocation for all vectors such that $c_{\overline{k}} \leq \frac{B}{\overline{k}} < c_{\overline{k}+1}$. As all cost distributions are identical, the probability to implement this payoff-maximizing set is maximized by setting these symmetric cutoffs.

The designer cannot additionally implement the first-best allocation for other cost vectors without violating at least one of the constraints. By setting cutoffs asymmetrically, the designer can greenlight $\overline{k}$ projects for other cost vectors. However, such an alternative mechanism features a lower probability to implement $\overline{k}$ projects.

Since for all other cost vectors the payoff-maximizing set is not implementable, the designer considers the next-best set, implementing the $(\overline{k} - 1)$ cheapest projects. Analogously to the steps before, she sets

$$z_i(\mathbf{c}_{-i}) = \min\{z^{**}, \frac{B}{\overline{k}-1}, c_{-i,\overline{k}-1}\} \quad \text{if } c_{-i,\overline{k}-1} \leq \frac{B}{\overline{k}-1}$$

to maximize the probability to implement the next-best set (greenlighting the cheapest $\overline{k}-1$ projects) taken as given the cutoffs set in Equations (11) and (12). We arrive at the proposed mechanism by continuing in this fashion. □

To sum up, in the symmetric case, the optimal allocation rule takes a simple form: The cheapest projects are greenlighted and the mechanism greenlights as many projects as the budget allows, while each procured project receives the same compensation. Any project that is redlighted prefers this allocation status over having to conduct the project with the associated compensation. It can be easily verified that this allocation rule indeed inherits all the properties we derived in the asymmetric two-project case. Thus, the optimal allocation is implementable with price clocks.

There are two rationales for greenlighted projects to get the same transfer. First, as shown in the proof of Proposition 3, this cutoff rule maximizes the probability of getting as many projects as possible. Strategyproofness prevents the budget from being shifted away from projects with low cost reports to projects with high costs as in the full-information allocation. Therefore offering equal cutoffs is the best the designer can do. Second, as seen in (3), the rewritten maximization problem of the designer, the expected utility of the designer is given by the sum of virtual surpluses of greenlighted projects. Therefore she wants to greenlight those projects with the highest virtual surpluses. That goal is consistent with offering equal cutoffs to greenlighted projects and excluding those with higher cost. In the optimal allocation, greenlighted projects have higher virtual surplus than those which are not greenlighted. The compatibility of the two goals - get as many projects as possible and get those with the highest virtual surpluses - is a special feature of the symmetric case, as he have demonstrated in Proposition 1.

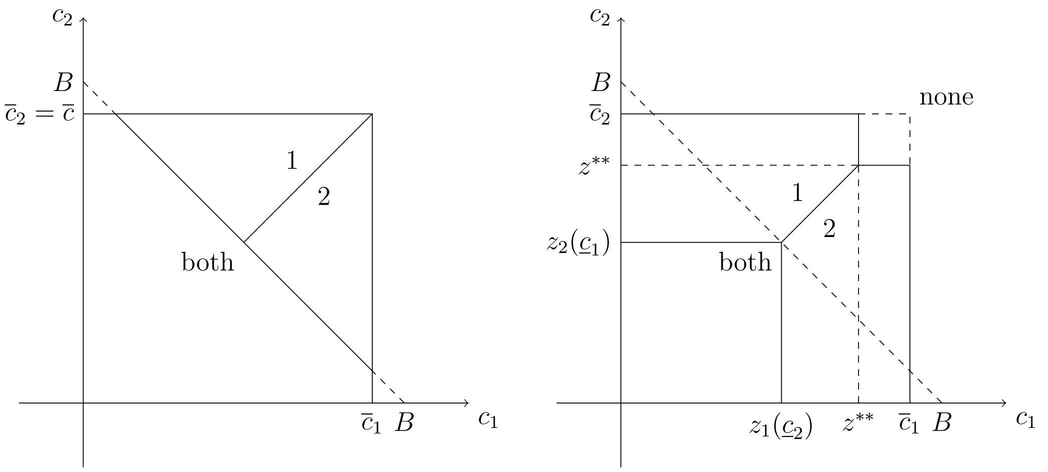

(a) Budget-constrained, full information. (b) Budget-constrained, private information.

Figure 5: An example of optimal allocations for the symmetric case with $n = 2$.

Figure 5 illustrates the optimal budget-constrained allocations in an example with two projects. Panel 5b shows the fully-constrained optimal allocation juxtaposed with the relaxed optimal allocation when (IC) is neglected, shown in Panel 5a. First, note that in this example $v \geq \overline{c}$ and $\overline{c} < B$. Therefore a fully-unconstrained designer with full information would always greenlight both projects, and a budget-constrained designer with full information would always greenlight at least one project. However, since $z^{**} < \overline{c}$, there exist realizations of $\mathbf{c}$ (the upper-right corner of Panel 5b) such that no project gets greenlighted in the (IC)-constrained optimal allocation, even though doing so would be profitable from an ex-post perspective. The negative virtual surpluses of the projects in these cases indicates that the cost of allocating to such a project - incentive compatibility requires higher transfers for other cost types - outweighs the benefit from an ex-ante perspective. The second major difference between the relaxed optimal allocation and the optimal allocation can be seen for those realizations of costs such that allocating to both projects would be feasible only in the relaxed problem. This difference is a result of the designer's inability to shift budget from low-cost to relatively higher-cost projects with a strategyproof mechanism.

Corollary 2. *In the symmetric case, the optimal direct mechanism can be implemented by a descending-clock auction. The clock price, denoted by τ, starts at z^{**} and descends continuously and synchronously down to $\frac{B}{n}$. Projects can drop out at any price but cannot re-enter. The auction stops once the clock price can be paid out to all projects remaining in the auction.*

In any iteration, a scoring function of the corresponding DA auction is

$$s_i^{A_t}(c_i, A_t) = \max\left\{c_i - \frac{B}{|A_t|}, 0\right\}.$$

We consider the descending-clock auction of Corollary 2 to be a natural indirect mechanism that implements the outcome of the optimal allocation. Project i's equilibrium strategy, which implements this outcome, has it staying active as long as the price is weakly larger than its private cost, $\tau \geq c_i$. It is easily verifiable that this is a weakly dominant strategy for project i.

3.4 $n > 2$. The asymmetric case

In general, not all our insights from the two-project case carry over as nicely as in the symmetric case. In this section, we first establish that the sufficient properties for implementability by clock auctions continue to hold if project substitutability holds, Lemma 5. However, it turns out that there are two major reasons why an optimal allocation rule may not have substitutes, as discussed with Example 2[14] and 3. Let us start with a generalized form of Lemma 3 assuming that project substitutability holds.

Lemma 5. *Suppose γ has substitutes. For any cost vectors $(\mathbf{c}_G, \mathbf{c}_R), (\mathbf{c}'_G, \mathbf{c}_R) \in C$ such that $G = \gamma(\mathbf{c}_G, \mathbf{c}_R) = \gamma(\mathbf{c}'_G, \mathbf{c}_R)$ and $R = I \setminus G$, the optimal cutoff*

[14]The spirit of this example is due to Daniel García.

function z_g for all $g \in G$ is (almost everywhere) independent of the costs of all greenlighted projects $\mathbf{c}_G$. That is,

$$z_g(\mathbf{c}_{G-g}, \mathbf{c}_R) = z_g(\mathbf{c}'_{G-g}, \mathbf{c}_R) \quad \forall g \in G.$$

The proof is stated in the appendix. It generalizes the intuition of the two-project proof by defining some a_i similar to (5), which is set-individual and contingent on cost reports of redlighted projects. The proof relies on the substitutes condition in Inequality (14). Unlike the two-project case, with more projects there not necessarily exists a cost combination $\dot{\mathbf{c}}$ such that $z_g(\dot{\mathbf{c}}_{-g}) = a_g^G$ for all projects g in some allocation set G. As an immediate consequence of Lemma 5, the following corollary establishes that any optimal cutoff mechanism with substitutes has non-bossy winners. That is, for optimal allocation rules with substitutes a clock-auction implementation exists.

Corollary 3. *Any optimal mechanism with substitutes also has non-bossy winners: If γ has substitutes, for all vectors $\widehat{\mathbf{c}}_G : \widehat{c}_g \leq \widetilde{c}_g$ for all $g \in G$,*

$$G = \gamma(\widetilde{\mathbf{c}}_G, \widetilde{\mathbf{c}}_R) \qquad \textit{implies} \qquad \gamma(\widehat{\mathbf{c}}_G, \widetilde{\mathbf{c}}_R) = G.$$

Hence, for all $i \in I$, for all $\mathbf{c}_{-i} \in C_{-i}$, and for all $\widehat{c}_i$, $\widetilde{c}_i \in C_i$ with $\widehat{c}_i < \widetilde{c}_i$, in any optimal mechanism,

$$\widehat{c}_i < \widetilde{c}_i \leq z_i(\mathbf{c}_{-i}) \qquad \textit{implies} \qquad \gamma(\widehat{c}_i, \mathbf{c}_{-i}) = \gamma(\widetilde{c}_i, \mathbf{c}_{-i}).$$

However, there are settings such that project substitutability is not optimal. Importantly, complementarities can ensue endogenously despite our assumption that projects' values and costs are independent of the allocation. If the lowest possible cost levels $\underline{c}_i$ are such that greenlighting some project combinations is never feasible, projects can endogenously become complements. We call such parameter combinations disjoint and also refer to two projects as disjoint if for no cost vector both projects can feasibly be greenlighted together. The first reason for complementarity is that two projects are only desirable when implemented together as seen in the following example: The designer prefers implementing 1 and 2 together over implementing 3 alone, but once either 1 or 2 becomes too expensive the other project is dropped as well in favor of implementing only project 3.

Example 2. *Suppose $n = 3$ and $z_i^{**} = \overline{c}_i$ for all $i \in \{1, 2, 3\}$. Let the values be such that*

$$\psi_1(c_1) + \psi_2(c_2) > \psi_3(c_3) > \max\{\psi_1(c_1), \psi_2(c_2)\}$$

for all $(c_1, c_2, c_3) \in C$ *and let the cost supports be such that*

$$\underline{c}_1 + \underline{c}_2 < \overline{c}_3 \leq B < \min\{\underline{c}_3 + \min\{\underline{c}_1, \underline{c}_2\}, \overline{c}_1, \overline{c}_2\}.$$

Then, the corresponding optimal mechanism has the following form

$$z_1(c_2, c_3) = \begin{cases} z \text{ if } c_2 \leq B - z \\ 0 \text{ otherwise} \end{cases}, z_2(c_1, c_3) = \begin{cases} B - z \text{ if } c_1 \leq z \\ 0 \text{ otherwise} \end{cases},$$

$$z_3(c_1, c_2) = \begin{cases} \overline{c}_3 \text{ if } c_2 > B - z \text{ or } c_1 > z \\ 0 \text{ otherwise} \end{cases}$$

This example can be seen as an extension of the asymmetric two-project case to a third project that can never be implemented together with any of the other two. To find the optimal allocation, the designer hast to find a constant z such that $\{1, 2\}$ is implemented when $(c_1, c_2) \leq (z, B - z)$ as in the two-project case. However, she does not have to consider the quantity-quality tradeoff. The reason is that project 3 is optimally greenlighted once one of the other projects' costs exceeds its cutoff.

Here, bidder substitutability fails because, as c_1 increases ($\varepsilon > 0$) from $z - \varepsilon$ to $z + \varepsilon$, project 2 with costs $c_2 \leq B - z$ gets dropped from the allocation set. The designer cannot consider an alternative mechanism that reduces z_3 marginally to increase z_2 because the lower cost bounds prohibit that projects 2 and 3 are ever conducted together and implementing $\{3\}$ is preferred to implementing $\{2\}$ alone. The cutoff mechanism in this example has non-bossy complements as projects 1 or 2 can only influence the allocation by changing their own allocation status.

This example satisfies only two of the three sufficient conditions for an implementation by a DA auction or clock auction. Clearly, the substitutes condition we imposed is not necessary. In this disjoint example, it is easy to construct an implementation with price clocks: All clocks start at the upper bounds. Then (at arbitrary speed) the prices of 1 and 2 descend to $(z, B - z)$. If both projects are still active, the price for project 3 jumps to zero, and 1 and 2 are implemented with their corresponding clock prices. If any project $i \in \{1, 2\}$ drops out earlier, then the price for $j \neq i$, $j \in \{1, 2\}$ drops to zero, while price 3 remains at $\overline{c}_3$: 3 is implemented with its maximal transfer.

Because project 1 and 2 are complements their price clocks have to be interconnected. This interconnection of price clocks is not a contradiction to the requirement that a DA scoring function of any project i only depends on c_i and the costs of rejected projects, $\mathbf{c}_{I\setminus A}$. The reason is that a single project $i \in \{1, 2\}$ can infer from its own cost whether allocation $\{1, 2\}$ is ruled out. Thus, either one of the two projects is rejected in the first iteration (has the highest score) or, if both projects find the payoff-maximizing allocation set to be feasible, project 3 is rejected and the algorithm stops in the next iteration. If a project $i \in \{1, 2\}$ is rejected in the first iteration, the score of the remaining project $j \in \{1, 2\}$ can depend on the cost of the rejected project i to be rejected next. Returning to a deferred acceptance logic, the scoring function first test whether the most-preferred allocation is blocked by project 1 or 2.

The next example features another kind of complementarity. In this example, project 3 can be a bossy loser. To construct this example, we intertwine a disjoint two-project symmetric environment with an additional small project. The small project can additionally be implemented if the residual budget suffices, which by construction of the example is not always the case.

Example 3. *Suppose $n = 3$ and $z_i^{**} = \overline{c}_i$ for all $i \in \{1, 2, 3\}$. Let*

$$\psi_1(c) = \psi_2(c) > \psi_3(c_3)$$

for all $(c, c, c_3) \in C$ and let $F_1 = F_2$ (implying $\underline{c}_1 = \underline{c}_2$) with cost supports such that

$$\min\{2\underline{c}_1, \underline{c}_1 + \overline{c}_3\} > B > \underline{c}_1 + \underline{c}_3.$$

The optimal mechanism takes the following form

$$z_1(c_2, c_3) = c_2, \quad z_2(c_1, c_3) = c_1, \quad z_3(c_1, c_2) = B - \max\{c_1, c_2\}.$$

Here, the designer prefers to implement the cheaper of the symmetric projects and adds small project 3 when feasible. The redlighted symmetric project can be a bossy loser since it determines the transfer to the other symmetric project. Hence, it determines the residual budget for project 3 and this residual budget is the cutoff level of project 3. As a result, an increase in the redlighted symmetric project's cost can, all else equal, kick project 3 out of the allocation set without affecting the rejected project's status. Again, there exists a straightforward DA-auction (clock-auction) implementation. Once a symmetric project is rejected, the second-iteration scoring function of project 3 can check whether the rejected project blocked the implementation of project 3.

In the above examples, the allocation rules are monotonic and have non-bossy winners. However, they only satisfy a weaker form of project substitutability: All projects are substitutes for disjoint projects, i.e., the cutoff of a project is weakly increasing in the cost of disjoint projects. For Example 2, we can readily verify that z_1 and z_2 are constant in c_3, and z_3 is weakly increasing in c_1 and c_2. Similarly, in Example 3, z_1 is increasing in c_2 and vice-versa. While z_1 is weakly decreasing in c_2 and vice-versa in Example 2 and while z_3 is weakly decreasing in c_1 and c_2 in Example 3, these projects are not disjoint with each other.

Weakening the substitutes condition is in the spirit of the matching literature where substitutability is known as a sufficient condition for many results, while it is clearly not always necessary in matching with contracts, see, e.g., Hatfield and Kojima (2008). While there is no less-restrictive sufficient conditions for all matching insights relying on substitutes, there are plausible weaker sufficient conditions for some known results. As a successful example, Hatfield and Kojima (2010) introduce the concept of bilateral substitutes as a sufficient condition for a stable matching to exist in the canonical doctors-match-hospitals model. It remains a question for future research to identify the weakest substitutes condition allowing for a DA-auction implementation in our setting.

4 Discussion

With our model as a starting point, there are several interesting modifications. In this section, we address the most natural alternative models or extensions.

$\mathbf{v}_i$ as private information, potentially correlated with $\mathbf{c}_i$ - The designer can neglect asking for v_i directly since no meaningful non-babbling equilibria in the v_i-dimension exist. If the conditional density of $v_i|c_i$ has full support, project i cannot credibly announce being a "high" type, say $\overline{v}_i$. If we slightly change the regularity assumption such that $\mathbb{E}[v_i|c_i] - c_i - \frac{F(c_i)}{f(c_i}$ must be strictly increasing, our results generalize by exchanging the previously commonly known v_i with $\mathbb{E}[v_i|c_i]$. This regularity condition mildly restricts the degree of positive correlation.

Interdependent types - We can interpret the symmetric case as a setting in which identical projects are provided at individual costs. Hence, one may wonder about a setting in which projects only draw an imperfect signal about the cost, which finally depends on other projects' signals as well. In a clock auction in such an environment, active projects update their belief about the cost whenever a project drops out. Moreover, the designer learns this information as well. Therefore the design of the optimal mechanism crucially depends on the information structure.

Residual money - Whether it is reasonable to assume that the designer values residual money depends on the application. In Ensthaler and Giebe (2014a), money does not enter the objective function, only the constraints. To clarify the relation to their paper, we introduce a linear weighting $\lambda \in [0,1]$ of residual money, and provide comparative statics on parameter λ. The objective function can be rewritten as in (3),

$$\max_{\{z_i\}_{i\in I}} \mathbb{E}_{\mathbf{c}} \left[\sum_i \mathbb{I}(c_i \leq z_i(\mathbf{c}_{-i})) \left(v_i - \lambda \left(c_i + \frac{F_i(c_i)}{f_i(c_i)} \right) \right) \right]$$
$$\text{s.t.}$$
$$\sum_{i\in I} \mathbb{I}(c_i \leq z_i(\mathbf{c}_{-i})) z_i(\mathbf{c}_{-i}) \leq B \quad \forall \mathbf{c} \in C.$$

This objective function highlights one difference to the original setting. Instead of $\boldsymbol{\zeta}^{**}$-exclusive the optimal mechanism is $\boldsymbol{\zeta}_\lambda^{**}$-exclusive: Define $\psi_{i,\lambda}(c) = v_i - \lambda(c + \frac{F_i(c)}{f_i(c)})$ as the λ-adjusted virtual surplus and define the vector $\boldsymbol{\zeta}_\lambda^{**}$ with i-the element $z_{i,\lambda}^{**} = \min\{\bar{c}_i, \psi_{i,\lambda}^{-1}(0)\}$.

Our insights in this paper qualitatively extend to any linear weighting λ. In fact, the optimal allocation in the symmetric case remains unchanged if $\boldsymbol{\zeta}_\lambda^{**} = (\bar{c}_1, \bar{c}_2, \ldots, \bar{c}_n)$ for all $\lambda \in [0,1]$, i.e., when the original optimal mechanism did not exclude any cost types. For any combination of cost supports and values, there exists a sufficiently small $\lambda' > 0$ such that the designer's ranking over projects is lexicographic. In other words, λ' must be sufficiently small such that no λ'-weighted difference in cost can offset any difference in values.

Introducing a weight λ affects the quantity-quality tradeoff. To illustrate how the optimal allocation varies when λ is perturbed, we consider Example 1 again, see Figure 6. A lower λ means that the designer prefers the high-value project 1 for higher cost reports relative to the low-value project 2 for a given cost report. This difference is illustrated by a right-shift in the diagonal that represents the loci such that both projects have equal (λ-adjusted) virtual surplus.

Reducing the weight of residual money increases the measure of cost reports for which the optimal mechanism implements project 2 despite project 1 having the larger λ-adjusted virtual surplus. As illustrated in Figure 6, reducing λ means that, in the optimal mechanism, the cutoffs at which both projects are greenlighted move southeast along the budget line, thus reducing the probability to greenlight both projects. The reason is that for lower λ a higher weight is placed on the high-value project 1.

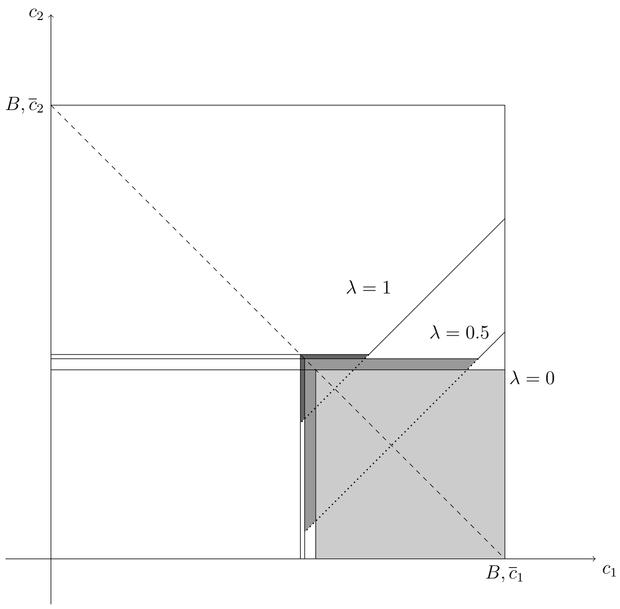

Figure 6: Decreasing λ augments the quantity-quality tradeoff: The gray areas, where the project with lower λ-adjusted virtual surplus is implemented, increases.

5 Conclusion

Despite their importance, knapsack problems with private information have been somewhat overlooked by the economics literature. We examine a setting in which a budget-constrained procurer faces privately-informed sellers under ex-post constraints. Amongst many possible economic problems, this setting particularly applies to development funds, which are typically endowed with a fixed budget and want to distribute this money to a set of heterogeneous projects. Such problems often entail relationships in which sellers can renege on the terms of the agreement ex-post. To avoid nondelivery, shelving the project or costly renegotiation, it is appropriate to impose ex-post constraints on the agents' participation. For a relevant subset such settings, we have shown that DA auctions constitutes the class of optimal deterministic strategyproof mechanisms.

An optimal mechanism is described by a set of cutoff functions: All projects that report costs below their cutoff are implemented and receive a transfer equal to their cutoff. In any two-project case, these cutoff functions are weakly increasing in the other project's costs, which means that the optimal allocation rule has substitutes: Given a project is implemented for some cost vector, it is also implemented when, all else being equal, the cost of the rival project is increased. For any optimal allocation rule that has substitutes, we show that it also has non-bossy winners: A project that is implemented cannot affect the allocation without changing its own allocation status. In particular, if two different realizations of the cost vector lead to the same allocation, then the cutoffs of conducted projects only vary in the costs of projects not conducted. Finally, the optimal allocation rule excludes all projects with negative "virtual surplus" from the allocation.

These properties allow for a characterization as a deferred acceptance (DA) auction, introduced by Milgrom and Segal (2015). The DA auction representation provides a simple implementation via descending-clock auctions, which are easy to understand and usable in practice. In addition, DA auctions have attractive properties regarding incentive compatibility which make the prediction of equilibrium play more robust. Furthermore, we investigate exemplary settings in which project substitutability fails, but a DA-auction implementation exists nevertheless. Thereby, we shed light on the necessity of substitute-like conditions.

We fully describe the optimal allocation and the corresponding descending-clock auction in an environment in which projects are ex-ante symmetric. The optimal mechanism is monotone in the sense that the cheapest projects are greenlighted and all projects conducted receive the same transfer. This transfer either corresponds to the lowest cost among redlighted projects or the budget is distributed equally. The equivalent clock auction features a single price clock that continuously descends until all active projects can be financed.

For asymmetric environments, in which values and/or cost distributions differ, we demonstrate a novel tradeoff between quantity and quality of the implemented projects. The designer prefers projects with high virtual surplus over projects with low virtual surplus and she prefers more projects over fewer projects. In models in which the designer wants to procure a fixed number of projects, she would always choose the projects with the highest virtual surpluses. If quantity is endogenously determined by a budget-constrained mechanism designer, it is ex-ante not always desirable to select the best projects. When the best projects are always conducted, incentive compatibility would force the designer to reduce the expected number of greenlighted projects. This insight entails a consequence for the corresponding descending-clock auction: Clocks not only run asynchronously, but also periodically have to stop for certain projects.

We identify an interesting question for future research, namely, what is the weakest substitute condition such that a DA implementation exists. Having an understanding of such a condition paves the way to study extensions such as multiple projects per agent and projects that are complements by assumption instead of (exogenous) substitutes. For practitioners, a simple approximately optimal detail-free mechanism may be of great value. The characterization of the optimal mechanism as a DA auction sheds light on how to construct such an approximately optimal mechanism such as Ensthaler and Giebe (2014b). Halting clocks should be a key feature for the corresponding clock auction in asymmetric environments. While we provide an elegant indirect mechanism, this mechanism is only easy to implement when details of the environments are known.

6 Appendix

Lemma 5. *Suppose γ has substitutes. For any cost vectors $(\mathbf{c}_G, \mathbf{c}_R), (\mathbf{c}'_G, \mathbf{c}_R) \in C$ such that $G = \gamma(\mathbf{c}_G, \mathbf{c}_R) = \gamma(\mathbf{c}'_G, \mathbf{c}_R)$ and $R = I \setminus G$, the optimal cutoff function z_g for all $g \in G$ is (almost everywhere) independent of the costs of all greenlighted projects $\mathbf{c}_G$. That is,*

$$z_g(\mathbf{c}_{G-g}, \mathbf{c}_R) = z_g(\mathbf{c}'_{G-g}, \mathbf{c}_R) \quad \forall g \in G.$$

Proof. When $\gamma(\mathbf{c})$ is a singleton, i.e., when only one project is greenlighted, the statement follows from the nature of a cutoff function. Take a feasible candidate mechanism with a set of (by substitutability) weakly increasing cutoff functions $\{z_i\}_{i \in I}$. Assume that given some cost vector $\widetilde{\mathbf{c}}_R$ there exists a set of cost vectors $\mathbf{c}_G$ with positive Lebesgue-measure such that for all those cost vectors $\gamma(\mathbf{c}_G, \widetilde{\mathbf{c}}_R) = G$. Then $a_i^G(\widetilde{\mathbf{c}}_R)$ according to the following definition

$$\begin{aligned} a_i^G(\widetilde{\mathbf{c}}_R) = \max\{c_i | \exists \mathbf{c}_{G-i} : \, & c_i \leq z_i(\mathbf{c}_{G-i}, \widetilde{\mathbf{c}}_R), \\ & \text{and } c_g \leq z_g(\mathbf{c}_{G-g}, \widetilde{\mathbf{c}}_R) \forall g \in G, \\ & \text{and } \widetilde{c}_r > z_r(\mathbf{c}_G, \widetilde{\mathbf{c}}_{R-r}) \forall r \in R\} \end{aligned} \tag{13}$$

exists for all $i \in G$.

In words, $a_i^G(\widetilde{\mathbf{c}}_R)$ is the highest cost of project i such that, given some cost vector $\widetilde{\mathbf{c}}_R$ of projects that are not executed, there exists some vector $\mathbf{c}_{G-i}$ of costs of competing projects that induces a cutoff $z_i(\mathbf{c}_{G-i}, \widetilde{\mathbf{c}}_R)$ above said cost while each element c_g of the vector $\mathbf{c}_{G-i}$ is lower than the cutoff induced by $a_i^G(\widetilde{\mathbf{c}}_R)$ and the elements of the cost vectors $\widetilde{\mathbf{c}}_R$ and $\mathbf{c}_{G-i-g}$,

$$\forall g \in G \setminus \{i\}, \;\; c_g \leq z_g(\widetilde{\mathbf{c}}_R, \mathbf{c}_{G-i-g}, a_i^G(\widetilde{\mathbf{c}}_R)).$$

Simultaneously, it must hold that these costs induce a cutoff such that no project $r \in R$ is conducted for this cost vector,

$$\forall r \in R, \;\; \widetilde{c}_r > z_r(\widetilde{\mathbf{c}}_{R-r}, \mathbf{c}_{G-i}, a_i^G(\widetilde{\mathbf{c}}_R)).$$

By left-continuity of the cutoff functions (see proof of Lemma 3), the limit is reached from below and there exists at least one cost vector $(\widetilde{\mathbf{c}}_R, \dot{\mathbf{c}}_{G-i}, a_i^G(\widetilde{\mathbf{c}}_R))$ such that G is the set of executed projects and $a_i^G(\widetilde{\mathbf{c}}_R) = z_i(\widetilde{\mathbf{c}}_R, \dot{\mathbf{c}}_{G-i})$ holds.

By construction,

$$\dot{c}_g \leq a_g^G(\widetilde{\mathbf{c}}_R) \, \forall g \in G \setminus \{i\},$$

because, given $\widetilde{\mathbf{c}}_R$, there cannot exist a cost vector where only all projects in G are executed and the cost of project g exceeds $a_g^G(\widetilde{\mathbf{c}}_R)$.

By the assumption of substitutes, z_i is weakly increasing. Hence

$$a_i^G(\widetilde{\mathbf{c}}_R) = z_i(\widetilde{\mathbf{c}}_R, \dot{\mathbf{c}}_{G-i}) \leq z_i(\mathbf{a}_{G-i}^G(\widetilde{\mathbf{c}}_R), \widetilde{\mathbf{c}}_R), \tag{14}$$

where $\mathbf{a}_{G-i}^G$ is the vector of all a_g^G defined according to (13) except a_i^G. The same logic also applies to all other projects $g \in G$.

Consequently,

$$G \subseteq \gamma(\mathbf{a}_G^G(\widetilde{\mathbf{c}}_R), \widetilde{\mathbf{c}}_R).$$

and thus the budget constraint requires that

$$\sum_{g \in G} z_g(\mathbf{a}_{G-g}^G(\widetilde{\mathbf{c}}_R), \widetilde{\mathbf{c}}_R) \leq B. \tag{15}$$

From here on, the proof in the main text applies. $\square$

References

Aggarwal, G., and J. D. Hartline (2006): “Knapsack auctions,” in *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pp. 1083–1092. Society for Industrial and Applied Mathematics.

Anari, N., G. Goel, and A. Nikzad (2014): “Mechanism Design for Crowdsourcing: An Optimal 1-1/e Competitive Budget-Feasible Mechanism for Large Markets,” in *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pp. 266–275. IEEE.

Baron, D. P., and R. Myerson (1982): “Regulating a monopolist with unknown costs,” *Econometrica*, pp. 911–930.

Bei, X., N. Chen, N. Gravin, and P. Lu (2012): “Budget feasible mechanism design: from prior-free to bayesian,” in *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pp. 449–458. ACM.

Börgers, T., and P. Norman (2009): “A note on budget balance under interim participation constraints: the case of independent types,” *Economic Theory*, 39(3), 477–489.

Caillaud, B., and J. Robert (2005): "Implementation of the revenue-maximizing auction by an ignorant seller," *Review of Economic Design*, 9(2), 127–143.

Che, Y.-K., and I. Gale (1998): "Standard auctions with financially constrained bidders," *The Review of Economic Studies*, 65(1), 1–21.

Chen, N., N. Gravin, and P. Lu (2011): "On the approximability of budget feasible mechanisms," in *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pp. 685–699. SIAM.

Chung, K.-S., and J. C. Ely (2002a): "Ex-post incentive compatible mechanism design," *Mimeo, Northwestern University.*

——— (2002b): "Mechanism Design with Budget Constraint," *Mimeo, Northwestern University.*

Dizdar, D., A. Gershkov, and B. Moldovanu (2011): "Revenue maximization in the dynamic knapsack problem," *Theoretical Economics*, 6(2), 157–184.

Dobzinski, S., C. H. Papadimitriou, and Y. Singer (2011): "Mechanisms for complement-free procurement," in *Proceedings of the 12th ACM conference on Electronic commerce*, pp. 273–282. ACM.

Dütting, P., V. Gkatzelis, and T. Roughgarden (2014): "The performance of deferred-acceptance auctions," in *Proceedings of the fifteenth ACM conference on Economics and computation*, pp. 187–204. ACM.

Ensthaler, L., and T. Giebe (2014a): "Bayesian optimal knapsack procurement," *European Journal of Operational Research*, 234(3), 774–779.

——— (2014b): "A dynamic auction for multi-object procurement under a hard budget constraint," *Research Policy*, 43(1), 179–189.

Gershkov, A., J. K. Goeree, A. Kushnir, B. Moldovanu, and X. Shi (2013): "On the equivalence of Bayesian and dominant strategy implementation," *Econometrica*, 81(1), 197–220.

Hatfield, J. W., and F. Kojima (2008): "Matching with contracts: Comment," *The American Economic Review*, 98(3), 1189–1194.

——— (2010): "Substitutes and stability for matching with contracts," *Journal of Economic theory*, 145(5), 1704–1723.

Kellerer, H., U. Pferschy, and D. Pisinger (2004): *Knapsack problems.* Springer Science & Business Media.

LEDYARD, J., D. PORTER, AND R. WESSEN (2000): “A market-based mechanism for allocating space shuttle secondary payload priority,” *Experimental Economics*, 2(3), 173–195.

LI, S. (2015): “Obviously Strategy-Proof Mechanisms,” *Available at SSRN 2560028.*

LUTON, R., AND R. P. MCAFEE (1986): “Sequential procurement auctions,” *Journal of Public Economics*, 31(2), 181–195.

MANELLI, A. M., AND D. R. VINCENT (2010): “Bayesian and Dominant-Strategy Implementation in the Independent Private-Values Model,” *Econometrica*, 78(6), 1905–1938.

MASKIN, E. (2002): “How to Reduce Greenhouse Gas Emissions: An Application of Auction Theory,” *Nancy L. Schwartz Memorial Lecture, Northwestern University.*

MILGROM, P., AND I. SEGAL (2015): “Deferred-acceptance auctions and radio spectrum reallocation,” in *Mimeo, Stanford University.*

MOOKHERJEE, D., AND S. REICHELSTEIN (1992): “Dominant strategy implementation of Bayesian incentive compatible allocation rules,” *Journal of Economic Theory*, 56(2), 378–399.

MOULIN, H. (1999): “Incremental cost sharing: Characterization by coalition strategy-proofness,” *Social Choice and Welfare*, 16(2), 279–320.

MOULIN, H., AND S. SHENKER (2001): “Strategyproof sharing of submodular costs:budget balance versus efficiency,” *Economic Theory*, 18(3), 511–533.

MU’ALEM, A., AND N. NISAN (2008): “Truthful approximation mechanisms for restricted combinatorial auctions,” *Games and Economic Behavior*, 64(2), 612–631.

MYERSON, R. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, pp. 58–73.

PAI, M. M., AND R. VOHRA (2014): “Optimal auctions with financially constrained buyers,” *Journal of Economic Theory*, 150, 383–425.

SINGER, Y. (2010): “Budget feasible mechanisms,” in *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pp. 765–774. IEEE.

STRAUSZ, R. (2003): “Deterministic mechanisms and the revelation principle,” *Economics Letters*, 79(3), 333–337.