

Algorithmic Mechanism Design with Investment*

Mohammad Akbarpour[†] Scott Duke Kominers[‡]
Kevin Michael Li[§] Shengwu Li[¶] Paul Milgrom^{||}

First posted: February 25, 2020

This version: May 12, 2023

Abstract

We study the investment incentives created by truthful mechanisms that allocate resources using approximation algorithms. Some approximation algorithms guarantee nearly 100% of the optimal welfare in the allocation problem but guarantee only 0% when accounting for investment incentives. An algorithm's allocative and investment guarantees coincide if and only if its *confirming negative externalities* are sufficiently small. We introduce fast approximation algorithms for the knapsack problem that have no confirming negative externalities and guarantees close to 100% for both allocation and investment.

Keywords: Combinatorial optimization, Knapsack problem, Investment, Auctions, Approximation, Algorithms

JEL classification: D44, D47, D82

*An extended abstract of this work appeared in the *Proceedings of the 22nd ACM Conference on Economics and Computation*. We thank Eric Tang for fantastic research assistance, and thank Moshe Babaioff, Ben Brooks, Peter Cramton, Michal Feldman, Matthew Gentzkow, Paul Goldsmith-Pinkham, Yannai Gonczarowski, Nima Haghpanah, Andy Haupt, John William Hatfield, Nicole Immorlica, Matthew Jackson, Emir Kamenica, Zi Yang Kang, Eric Maskin, Ellen Muir, Rad Niazadeh, Noam Nisan, Amin Saberi, Roberto Saitto, Mitchell Watt, numerous seminar audiences, the editor (Bart Lipman), and several referees for helpful comments. We thank Broadsheet Cafe for inspiration and coffee. Akbarpour and Kominers gratefully acknowledge the support of the Washington Center for Equitable Growth. Additionally, Kominers gratefully acknowledges the support the National Science Foundation (grant SES-1459912) and both the Ng Fund and the Mathematics in Economics Research Fund of the Harvard Center of Mathematical Sciences and Applications. Milgrom gratefully acknowledges support from the National Science Foundation (grant SES-1947514). All errors remain our own.

[†]Stanford University. Email: mohamwad@stanford.edu

[‡]Harvard University and a16z crypto. Email: kominers@fas.harvard.edu

[§]Stanford University. Email: kli317@stanford.edu

[¶]Harvard University. Email: shengwu_li@fas.harvard.edu

^{||}Stanford University and Auctionomics. Email: milgrom@stanford.edu

1 Introduction

Many real-world allocation problems are too complex for exact optimization. For example, it is computationally difficult—even under full information—to optimally pack indivisible cargo for transport (Dantzig, 1957; Karp, 1972), to coordinate electricity generation and transmission (Lavai and Low, 2011; Bienstock and Verma, 2019), to assign radio spectrum broadcast rights subject to legally-mandated interference constraints (Leyton-Brown et al., 2017), or to find welfare-maximizing allocations in combinatorial auctions (Sandholm, 2002; Lehmann et al., 2006). Rising to the challenge for these and many other problems, researchers have developed fast approximation algorithms.

Approximation algorithms can be combined with pricing rules to produce truthful mechanisms, provided that the algorithm is “monotone” (Lavi et al., 2003). In this paper, we study the *ex ante* investment incentives created by such mechanisms.

Suppose that one bidder can make a costly investment to change its value before participating in a truthful mechanism. As an initial result, we show that all truthful mechanisms using the same allocation algorithm entail the same investment incentives, so we can regard the investment incentives as properties of the algorithm itself.

If the allocation algorithm exactly maximizes total welfare, then the corresponding truthful mechanism is a Vickrey-Clarke-Groves (VCG) mechanism. For VCG mechanisms, any single bidder’s investment is profitable if and only if it improves total welfare (Rogerson, 1992). In this respect, the VCG mechanisms are essentially unique. We find that a truthful mechanism aligns a bidder’s investment incentives with welfare maximization only if there is some set of allocations such that, for generic valuation profiles, its allocation algorithm exactly maximizes welfare over that set. Many practical approximation algorithms do not have this structure and, as a result, lack efficient investment incentives.

One might also hope that if an allocation algorithm approximately maximizes total welfare, then it generates approximately efficient investment incentives—but we show to the contrary that arbitrarily good approximations can have arbitrarily bad investment guarantees. To make this statement precise, we evaluate an algorithm’s performance on any particular instance by the welfare it achieves divided by the maximum welfare. We refer to the worst-case ratio over all instances when values are exogenous as the *allocative guarantee*, and the worst-case ratio when one bidder’s *ex ante* investment endogenously determines its value as the *investment guarantee*.¹ (The investment guarantee measures welfare net of investment costs.) Because the investment guarantee is a worst-case over instances and over

¹Our results partially extend to the case of multiple bidders who make simultaneous investments, as we discuss in Section 2.4.6.

investment technologies, it is never more than the allocative guarantee. We characterize the algorithms for which the allocative and investment guarantees are equal, and apply those results to evaluate and improve upon standard approximation algorithms.

1.1 The knapsack example

We use the knapsack problem (Dantzig, 1957) to introduce the investment problem and our general results. An instance of the knapsack problem is described by a list of indivisible items, each having a positive size and value, and a capacity constraint. Each item’s size is no more than the capacity constraint. The problem is to select (“pack”) a set of items to maximize the total value, subject to the sum of the item sizes not exceeding the capacity constraint. Finding an exact optimal solution to the knapsack problem is computationally difficult—it is NP-hard.

Suppose that each item is associated with a different bidder and that all the item sizes are publicly observed, but the value of being packed is the bidder’s private information. An algorithm for the knapsack problem is *monotone* if when any packed bidder’s value is increased, the algorithm still selects that bidder to be packed. Any monotone algorithm can be paired with a payment rule to create a truthful mechanism. One such payment rule—the *threshold rule*—charges zero to each unpacked bidder and charges each packed bidder its *threshold price*, which is the infimum of the set of values that would result in the bidder being packed.

A packing algorithm that is careless about packing low-value items can have a good allocative guarantee but a poor investment guarantee. For an extreme but simple example, consider a *satisficing algorithm* that packs only the most valuable item when its value is at least 99% of the sum of all values, and otherwise optimizes exactly. This algorithm is monotone, so its threshold auction is truthful, and its allocative guarantee is 99%. Its investment guarantee, however, is 0, as shown by the example in Table 1. In this example, suppose that bidder A can invest at a cost of 200 to raise its value from ε to $200 + 2\varepsilon$. This investment is profitable and causes the satisficing algorithm to pack just A , for net welfare of $200 + 2\varepsilon - 200 = 2\varepsilon$. But the social optimum is to invest and pack both A and B , for net welfare $1 + 2\varepsilon$. Thus, despite the 99% allocative guarantee, the investment guarantee of this algorithm is no more than $\inf_{\varepsilon > 0} \left\{ \frac{2\varepsilon}{1+2\varepsilon} \right\} = 0$.

While the preceding example is extreme, it contains the seed of a general lesson. In the example, increasing the value of a packed item worsens the packing of other items. We show that a monotone algorithm’s investment performance can be worse than its allocative performance *only if* an investment that “confirms” the investor’s allocation can reduce

bidder	A	B	C
value	ε	1	.5
size	.4	.6	.5

Table 1: A knapsack instance, with capacity 1. Assume $0 < \varepsilon < .2$.

the total value of other participants’ allocations—an effect that we call a *confirming negative externality*. An algorithm that excludes confirming negative externalities—an *XCONE algorithm*—always has an investment guarantee equal to its allocative guarantee.

1.2 Summary of main results

For our general treatment, we assume a finite set of outcomes. Each bidder’s value is a vector v_n , with element $v_{n,o}$ capturing bidder n ’s value for outcome o . The bidder’s possible values are a product of intervals, one for each outcome. An allocation assigns one outcome to each bidder; there is an arbitrary set of feasible allocations. An algorithm selects an allocation given the value profile v and the set of feasible allocations.

We allow that investments may be made under uncertainty about all the inputs to the algorithm, including the values resulting from the bidder’s investment, the values reported by the other bidders, and the set of feasible allocations. Each input is a random variable described by a function from a finite state space S . We allow any probability distribution on S ; our only restriction is that in every state, the realization is an allowable instance of the deterministic problem. We assume that the investor selects an investment from some finite set to maximize the investor’s own expected payoff, net of its investment cost and the price it pays in the truthful mechanism. We compare the resulting expected welfare to that of the *optimum*, which results from the *ex ante* efficient investment and the *ex post* efficient allocation. An algorithm is a β -*approximation for investment* if for any instance of the investment problem, the expected welfare of the mechanism is at least β times the optimum welfare.

We prove that the worst case is always deterministic, so introducing uncertainty does not affect the investment guarantee: if x is a β -approximation for investment for all singleton state spaces, then it is also a β -approximation for investment for any finite state space.

To study algorithm performance under certainty, we introduce two new concepts, *algorithmic externalities* and *confirming changes*.

When a bidder invests under a truthful mechanism, that bidder changes its reported value. We define the *algorithmic externality* from that change to be the increase (positive, negative or zero) in the sum of other bidders’ values for the resulting allocation, plus the increase in the price the bidder pays. For example, if the bidder’s change in value reduces

the welfare of other bidders but also increases the bidder’s payment to the auctioneer by a larger absolute amount, we count that as a positive externality because it increases the total welfare of the other bidders plus the auctioneer. As we show, any two truthful mechanisms that use the same underlying allocation algorithm result in identical externalities.

Suppose that at some value profile (v_n, v_{-n}) , the algorithm allocates outcome o to bidder n . A change in bidder n ’s values from v_n to \tilde{v}_n is *confirming* if for any other outcome \hat{o} , we have $\tilde{v}_{n,o} - v_{n,o} \geq \tilde{v}_{n,\hat{o}} - v_{n,\hat{o}}$; that is, n ’s value for the original outcome increases at least as much as its value for any other outcome. If the inequalities were all strict, then monotonicity of the algorithm would imply that such a value change must leave n ’s outcome unchanged, but others’ outcomes may change. With weak inequalities, it is possible that n ’s outcome changes as well, with a compensating change to n ’s payments. If given a confirming change to n ’s value, the algorithm’s allocation changes in a way that results in a negative externality, we call that a *confirming negative externality*.

Our first main result establishes a necessary and sufficient condition for the investment and allocative guarantees to coincide, in the form of a bound on the magnitude of confirming negative externalities. Suppose we start at value profile (v_n, v_{-n}) and make a confirming change to \tilde{v}_n ; our condition requires that any resulting negative externality must not exceed the slack in the allocative guarantee β^* at the original value profile. This bound can be hard to assess, however, because the slack depends on the optimal welfare, which is hard to compute or characterize for many problems of interest. The second result is a corollary of the first that may be more useful because it is easier to check: if an algorithm excludes confirming negative externalities (XCONE), then its investment and allocative guarantees coincide.

To explore the intuition for our results, we limit attention here to packing problems, in which bidder n faces only two outcomes—winning (being “packed”) or losing—when its price for being packed is p .² A preliminary observation is that the worst-case investment performance must occur when a bidder who chooses to invest c makes zero additional profit, because reducing c leaves the investment decision unchanged while increasing the ratio of the algorithm’s welfare net of costs to the optimal net welfare.

Consider a bidder n who would not be packed without investment, but by investing at cost $c > 0$ can increase its value to $p + c$, resulting in a net profit of zero. The intuitive argument hinges on decomposing this investment into two parts. Suppose that the bidder first has the option of investing at zero cost to raise its value just to the threshold price p . This is a zero-profit investment, since the result is that the bidder is packed but pays its

²Under the assumptions listed below, the same arguments and intuition can be extended to the case with multiple outcomes.

full value p for that. Because the investment cost is zero, this results in the same welfare as if the bidder’s value were fixed at p , so the allocative guarantee implies that welfare is at least β^* times the optimum. In the actual problem, bidder n can invest $c > 0$ to increase its value above the threshold to $p + c$. Compared to raising value to the threshold, adding this option changes neither the investor’s net payoff nor the total net welfare under the optimal benchmark. Thus, this investment, which confirms n ’s packing, reduces the algorithm’s performance for that instance if and only if it leads to a confirming negative externality. If that externality is large enough, it may drag total welfare below β^* times the optimum. For an XCONE algorithm, there are no confirming negative externalities, so there is no instance in which total net welfare is less than when the investor’s value is equal to the threshold, which is at least β^* times the optimal welfare for that problem. So, for an XCONE algorithm, investment performance is as good as allocation performance.

Some familiar approximation algorithms are XCONE; these include greedy algorithms for the knapsack problem and the clock auction algorithm used for the 2016 Federal Communication Commission’s broadcast incentive auction (with a perfect feasibility checker) (Leyton-Brown et al. (2017), Milgrom and Segal (2020)).

XCONE is also closely related to non-bossiness. An algorithm is non-bossy if a change in one bidder’s report that does not affect that bidder’s outcome cannot change any other bidder’s outcome. For allocation problems with two outcomes, such as the knapsack problem, every non-bossy algorithm is XCONE. Moreover, we show that any algorithm that takes the best output from a family of non-bossy algorithms is XCONE.³

We use our results to illuminate and repair the bad investment performance of a certain “fully polynomial time approximation scheme” (FPTAS), that is, an indexed collection of algorithms with parameter $\varepsilon > 0$ that have allocative guarantees of $1 - \varepsilon$ and maximum run-times bounded by a polynomial function of ε^{-1} and the input length. The Briest et al. (2005) (henceforth BKV) FPTAS for the knapsack problem consists of algorithms that are monotone but not XCONE. Each of the BKV algorithms fails XCONE because, like the satisficing algorithm described earlier, they pack low value items poorly when one item has a very high value. Just as in our earlier example, this results in the BKV algorithms having investment guarantees of zero. We show how to modify the BKV algorithms to be XCONE—creating a FPTAS for which the allocative and investment guarantees coincide.

³This result extends to problems with more than two outcomes under a tie-breaking condition.

1.3 Related work

Economists have studied *ex ante* investment in mechanism design at least since the work of [Rogerson \(1992\)](#), who demonstrated that [Vickrey](#) mechanisms induce efficient investment. [Bergemann and Välimäki \(2002\)](#) extended [Rogerson](#)'s finding in a setting with uncertainty, in which bidders invest in information before participating in an auction. Relatedly, [Arozamena and Cantillon \(2004\)](#), studied pre-market investment in procurement auctions, showing that while second-price auctions induce efficient investment, first-price auctions do not. [Hatfield et al. \(2014, 2019\)](#) extended these findings to characterize a relationship between the degree to which a mechanism fails to be truthful and/or efficient and the degree to which it fails to induce efficient investment. While like us, [Hatfield et al. \(2014, 2019\)](#) dealt with the connection between (near-)efficiency at the allocation stage and (near-)efficiency at the investment stage, they used additive error bounds, rather than the multiplicative worst-case bounds that are standard for the analysis of computationally hard problems. [Gershkov et al. \(2021\)](#) studied the construction of revenue-maximizing mechanisms with *ex ante* investment. [Tomoeda \(2019\)](#) studied full implementation of exactly-efficient social choice rules with investment.

Our paper is not the first to study investment incentives in an NP-hard allocation setting. [Milgrom \(2017\)](#) introduced a “knapsack problem with investment” in which the items to be packed are owned by individuals, and owners may invest to make their items either more valuable or smaller (and thus easier to fit into the knapsack). In the present paper, we reformulate the investment question in terms of worst-case guarantees and broaden the formulation to study truthful mechanisms for a wide class of resource allocation problems.

Our paper is naturally connected to the large literature on algorithmic mechanism design, started by the seminal paper of [Nisan and Ronen \(1999\)](#). This literature considers computational complexity in mechanism design, and explores properties of approximately optimal mechanisms. Among these works are those of [Nisan and Ronen \(2007\)](#) and [Lehmann et al. \(2002\)](#). [Nisan and Ronen \(2007\)](#) showed that VCG-based mechanisms with nearly-optimal allocation algorithms are generically non-truthful, while [Lehmann et al. \(2002\)](#) introduced a truthful mechanism for the knapsack problem in which the allocation is determined by a greedy algorithm. In addition, [Hartline and Lucier \(2015\)](#) developed a method for converting a (non-optimal) algorithm for optimization into a Bayesian incentive compatible mechanism with weakly higher social welfare or revenue; [Dughmi et al. \(2017\)](#) generalized this result to multidimensional types. For a more comprehensive review of results on approximation in mechanism design, see [Hartline \(2016\)](#).

There is also a literature on greedy algorithms, all of which are XCONE, which sort bidders based on some criterion and choose them for packing in an irreversible way; see [Pardalos et al. \(2013\)](#) for a review. [Lehmann et al. \(2002\)](#) studied the problem of constructing truthful

mechanisms from greedy algorithms; similarly, [Bikhchandani et al. \(2011\)](#) and [Milgrom and Segal \(2020\)](#) proposed clock auction implementations of greedy allocation algorithms.

Finally, our concept of an XCONE algorithm is closely related to the definition of a “bitonic” algorithm, introduced by [Mu’Alem and Nisan \(2008\)](#) to construct truthful mechanisms in combinatorial auctions. Bitonicity is defined for binary outcomes; with the restriction to binary outcomes, every XCONE algorithm is bitonic, but not vice versa.

2 The model and results

2.1 Approximation algorithms

Consider a set of **bidders** N and a set of **outcomes** O , both finite. For instance, in the knapsack problem, the set of outcomes is {packed, unpacked}. The value of bidder n for outcome o is $v_{n,o} \in \mathbb{R}_{\geq 0}$. We write $v_n \equiv (v_{n,o})_{o \in O}$ to denote n ’s **values**, and we write $v \equiv (v_n)_{n \in N}$ to denote a full **value profile**. An **allocation** $a \in O^N$ assigns one outcome to each bidder; a_n denotes the outcome of bidder n .

An **allocation instance** (N, O, v, A) consists of a set of bidders N , a set of outcomes O , a value profile v , and a set of **feasible allocations** $A \subseteq O^N$. To simplify notation, we often write instances as a pair (v, A) , leaving N and O implicit.

The standard approach in computer science is to assess an algorithm’s worst-case performance over a domain of instances. Hence, we define an **allocation problem** Ω to be a collection of instances.

Assumption 2.1. We assume that the value profiles in Ω have a product structure. That is, let $V_{n,o}$ be a closed interval of $\mathbb{R}_{\geq 0}$ capturing the possible values that bidder n might have for outcome o . We define $V_n \equiv \prod_{o \in O} V_{n,o}$ and require that $\{v : (N, O, v, A) \in \Omega\} = \prod_{n \in N} V_n$.

In some settings, one outcome $o \in O$ is an **outside option** known to be valued at 0; we capture this with $V_{n,o} = \{0\}$.

Assumption 2.1 is restrictive. For instance, suppose that each outcome is a bundle of goods. If a bidder has additive valuations, then their value for a bundle is equal to the sum of their values for the individual goods. Thus, the allocation problem comprised of all and only the instances with additive valuations does not satisfy Assumption 2.1.

An allocation algorithm is a computational procedure that takes as input the value profile and the set of feasible allocations and then outputs an allocation. Each algorithm induces an allocation rule, that is, a function from inputs to outputs. Practical algorithms must run quickly, but most of our results do not depend on running time, so we often use the

term “algorithm” to refer both to the computational procedure and to the function that it induces.

We restrict attention to deterministic allocation algorithms. Formally, an **algorithm** x is a function that selects a feasible allocation for each instance $(v, A) \in \Omega$; that is, $x(v, A) \in A$.⁴ We denote the outcome assigned to bidder n under x by $x_n(v, A)$. We abuse notation and identify outcomes $o \in O$ with binary vectors of length $|O|$, with one element equal to 1 and all others equal to 0, which allows us to write the **welfare** of algorithm x on instance (v, A) as

$$W_x(v, A) \equiv \sum_n [v_n \cdot x_n(v, A)].$$

The **optimal welfare** at instance (v, A) is

$$W^*(v, A) \equiv \max_{a \in A} \left\{ \sum_n [v_n \cdot a_n] \right\}.$$

Given some $\beta \in [0, 1]$, algorithm x is a **β -approximation for allocation** if for all $(v, A) \in \Omega$, we have that $W_x(v, A) \geq \beta W^*(v, A)$. We refer to the largest such β as the algorithm’s **allocative guarantee**.

2.2 Truthful mechanisms

Suppose that the bidders’ values are private information, so that the algorithm cannot directly input each bidder n ’s value v_n but must instead rely on each bidder’s *reported* value \hat{v}_n . To elicit these reports, we use a **mechanism** (x, p) , which is a pair consisting of an algorithm x and a payment rule p that maps any reported instance (\hat{v}, A) into an allocation $x(\hat{v}, A) \in A$ and a profile of payments $p(\hat{v}, A) \in \mathbb{R}^N$. We adopt the sign convention that payments are made by the participants and to the auctioneer. A mechanism is **truthful** if for all instances $(v, A) \in \Omega$ and all $\hat{v}_n \in V_n$, we have that

$$v_n \cdot x_n(v, A) - p_n(v, A) \geq v_n \cdot x_n(\hat{v}_n, v_{-n}, A) - p_n(\hat{v}_n, v_{-n}, A).$$

When can an algorithm be paired with a pricing rule to produce a truthful mechanism?

⁴In complexity theory, we often are not given the feasible allocations A directly, but instead only a description that implies which allocations are feasible. For instance, a description could specify the bidders’ sizes and the capacity of the knapsack. In principle, algorithms for the knapsack problem could output different allocations for two instances with different item sizes but the same feasible allocations. Our formulation ignores this description-dependence, but we could easily accommodate it by specifying a function from descriptions to feasible allocations, and defining an instance as consisting of a value profile v and a description d ; none of our results would materially change with this adjustment.

Algorithm x is **weakly monotone (W-Mon)** if for any two instances (v, A) and (\tilde{v}_n, v_{-n}, A) , we have

$$[\tilde{v}_n - v_n] \cdot [x_n(\tilde{v}_n, v_{-n}, A) - x_n(v, A)] \geq 0. \quad (1)$$

For **packing problems**, we have $O = \{\text{packed}, \text{unpacked}\}$, a value for being “packed” $v_{n,\text{packed}} \geq 0$, and an outside option with $v_{n,\text{unpacked}} = 0$. In this special case, (1) reduces to the requirement that, under x , if n is packed at (v_n, v_{-n}, A) and $\tilde{v}_{n,\text{packed}} \geq v_{n,\text{packed}}$, then n is packed at (\tilde{v}_n, v_{-n}, A) .

A necessary condition for the existence of a pricing rule p such that (x, p) is truthful is that the algorithm is weakly monotone—and since V_n is convex by Assumption 2.1, this is also sufficient.⁵

Lemma 2.2 (Lavi et al. (2003); Saks and Yu (2005)). *An algorithm x is weakly monotone if and only if there exists a payment rule p such that (x, p) is truthful.*

Pricing rules in truthful mechanisms can be defined in terms of threshold prices, one for each outcome. The least value for bidder n to achieve outcome o is denoted by

$$\hat{\tau}_{n,o}(v_{-n}, A, x) = \inf_{v_n \in V_n} \{v_{n,o} : x_n(v_n, v_{-n}, A) = o\}$$

and the **threshold price** is

$$\tau_{n,o}(v_{-n}, A, x) \equiv \min \{\hat{\tau}_{n,o}(v_{-n}, A, x), \sup V_{n,o}\}.$$

Our results hold trivially if $\tau_{n,o}(v_{-n}, A, x) = \infty$ for some o .⁶ To focus on the non-trivial case, we assume that $\tau_{n,o}(v_{-n}, A, x) < \infty$; since $V_{n,o}$ is closed, it then follows that $\tau_{n,o}(v_{-n}, A, x) \in V_{n,o}$. We denote the **threshold vector** by $\tau_n(v_{-n}, A, x) \equiv (\tau_{n,o}(v_{-n}, A, x))_{o \in O}$. The set of possible values V_n has a product structure, so we have $\tau_n(v_{-n}, A, x) \in V_n$.

Now, V_n is path-connected, so a standard argument using the envelope theorem yields the following lemma (Milgrom and Segal, 2002).

Lemma 2.3. *If (x, p) is a truthful mechanism, then for each n , there exists a real-valued function $f_n(v_{-n}, A)$ such that*

$$p_n(v, A) = \tau_n(v_{-n}, A, x) \cdot x_n(v, A) + f_n(v_{-n}, A).$$

⁵Bikhchandani et al. (2006) provided other domain assumptions such that weak monotonicity is sufficient.

⁶Observe that if $\tau_{n,o}(v_{-n}, A, x) = \infty$, then $\hat{\tau}_{n,o}(v_{-n}, A, x) = \infty$ and $\sup V_{n,o} = \infty$, i.e., bidder n is never allocated outcome o and can have arbitrarily large values for o , which in turn implies that x has an allocative guarantee of 0.

Lemma 2.3 states that in a truthful mechanism, each bidder pays the threshold price to achieve its assigned outcome plus a strategically irrelevant term that does not depend on the bidder’s own report. Truthfulness of (x, p) implies that x assigns each bidder an outcome that maximizes its value minus its threshold price.

2.3 Algorithmic externalities

Given mechanism (x, p) and instance (v, A) , the **externality** of changing n ’s value from v_n to \tilde{v}_n is

$$\mathcal{E}_{x,p}(\tilde{v}_n, (v, A)) \equiv \underbrace{p_n(\tilde{v}_n, v_{-n}, A) - p_n(v, A)}_{\text{change in } n\text{'s payment}} + \underbrace{\sum_{m \neq n} v_m \cdot [x_m(\tilde{v}_n, v_{-n}, A) - x_m(v, A)]}_{\text{effect on others' welfare}}. \quad (2)$$

Expression (2) is the portion of n ’s effect on other participants’ welfare that is not fully reflected by n ’s price.⁷ Equivalently, if we treat the auctioneer as the residual claimant to any surplus or deficit of the mechanism, then (2) is the change in the sum of the payoffs of other participants, including the auctioneer.

Lemma 2.3 implies that any two truthful mechanisms that use the same allocation algorithm x have the same externalities. Consequently, we henceforth suppress the dependence of $\mathcal{E}_{x,p}$ on p , writing \mathcal{E}_x and calling this an **algorithmic externality**. VCG mechanisms have zero externalities, so it follows that if x is exactly maximizing, then x has no algorithmic externalities.

For an algorithm to yield efficient investment incentives, it must have zero externalities. Suppose that bidder n changes its value from v_n to \tilde{v}_n . If $\mathcal{E}_x(\tilde{v}_n, (v, A)) \neq 0$, then the change in n ’s payment does not fully capture the effect on others’ welfare. Thus, we can find cost $c \in \mathbb{R}$ such that paying c to change n ’s value from v_n to \tilde{v}_n is privately profitable but not socially optimal, or socially optimal but not privately profitable.

We characterize the zero-externality algorithms. An algorithm x is **maximal-in-range** if for each set of feasible allocations A , there exists an $R \subseteq A$ such that for all v we have

$$x(v, A) \in \operatorname{argmax}_{a \in R} \left\{ \sum_n [v_n \cdot a_n] \right\}.$$

We say that algorithm x is **welfare-equivalent** to algorithm x' if for every instance (v, A) we have $W_x(v, A) = W_{x'}(v, A)$. Note that if two algorithms are welfare-equivalent, then they

⁷In some parts of the economics and mechanism design literatures, the word “externality” is used to refer just to the second term, but our definition here is faithful to the traditional Pigouvian concept of externality.

yield identical allocations except when two allocations yield exactly the same welfare.

Theorem 2.4. *Suppose that algorithm x is weakly monotone. Then x has no algorithmic externalities ($\mathcal{E}_x \equiv 0$) if and only if x is welfare-equivalent to a maximal-in-range algorithm.*

Theorem 2.4 implies that it is (essentially) only VCG mechanisms that have no externalities, since any truthful mechanism based on a maximal-in-range algorithm is just a VCG mechanism with restricted range.⁸ Some allocation problems have fast maximal-in-range algorithms with meaningful allocative guarantees (Holzman et al., 2004; Dobzinski and Nisan, 2007).

Theorem 2.4 is substantially the same as Theorem 3.2 of Nisan and Ronen (2007), except that the Nisan and Ronen (2007) version requires the possible values V_n to be unbounded above, whereas ours allows V_n to be any product of closed intervals.

2.4 Performance under investment

In mechanisms with algorithmic externalities, selfish investment decisions do not always maximize social welfare. Thus, we study the connection between algorithmic externalities and performance guarantees under investment.

Given a truthful mechanism (x, p) , we assess whether the mechanism’s allocative guarantee also applies to investment problems in which a single bidder, denoted $\iota \in N$, can decide *ex ante* whether to invest and/or what investment to make. In our formulation, the bidder may be uncertain about what the situation will be when the mechanism is run, including potential uncertainty as to the values that would result from each of its possible investments, the values of the other bidders, and the feasible set that will apply. We compare the expected social welfare from the bidder’s selfish investment choice and the given mechanism’s allocation to the expected welfare from making the *ex ante* efficient investment and using the *ex post* efficient allocation.

We model the investor’s uncertainty using a probability space with a finite number of states S . Each uncertain **investment** opportunity is a pair (ν_ι, c) consisting of a function $\nu_\iota : S \rightarrow V_\iota$ and a cost $c \in \mathbb{R}$. An **investment instance** specifies the set of possible investments as well as the other bidders’ values as a function $\nu_{-\iota} : S \rightarrow V_{-\iota}$ and the feasible set as a correspondence $\mathcal{A} : S \rightrightarrows A$.⁹

⁸Here is an example of a weakly monotone, zero-externality algorithm that is not maximal-in-range, which applies to the problem of selecting two auction winners from among four bidders: If all four bidders have the same value of winning, then the algorithm selects bidders 1 and 2; otherwise, it selects a pair of bidders to maximize welfare from the set Φ consisting of the other five bidder pairs. This algorithm is welfare-equivalent to the algorithm that always selects the welfare-maximizing allocation from Φ .

⁹We do not restrict the correlations among these uncertain elements.

Formally, an investment instance $(N, O, S, g, \iota, I, \nu_{-\iota}, \mathcal{A})$ consists of:

1. Sets of bidders N and outcomes O .
2. A finite set of states S and a probability distribution $g \in \Delta S$.
3. A distinguished bidder—the **investor**—which we denote by the Greek letter ι .
4. A finite set of investments I for ι . To represent the status quo, we require that this set includes at least one pair (ν_ι, c) with $c = 0$.¹⁰
5. A function from states to the other bidders' values, $\nu_{-\iota} : S \rightarrow V_{-\iota}$.
6. A correspondence from states to feasible allocations, $\mathcal{A} : S \rightrightarrows O^N$.

We require that each state $s \in S$ and investment $(\nu_\iota, c) \in I$ together result in an instance of the original allocation problem, i.e., that $(N, O, \nu(s), \mathcal{A}(s)) \in \Omega$. To simplify our notation, we write each investment instance in the form $\bar{\omega} = (g, I, \nu_{-\iota}, \mathcal{A})$, suppressing N, O, S , and ι . (We use an overline to distinguish functions or variables related to an investment problem.)

Suppose that the investor participates in some truthful mechanism (x, p) . After an investment is chosen and the state is realized, the investor can do no better than to report the resulting value to the mechanism truthfully. Hence, its best response investment choice at instance $\bar{\omega} = (g, I, \nu_{-\iota}, \mathcal{A})$ is

$$\text{BR}(x, p, \bar{\omega}) \equiv \operatorname{argmax}_{(\nu_\iota, c) \in I} \left\{ \left(\sum_{s \in S} g(s) [\nu_\iota(s) \cdot x_\iota(\nu_\iota(s), \nu_{-\iota}(s), \mathcal{A}(s)) - p_\iota(\nu_\iota(s), \nu_{-\iota}(s), \mathcal{A}(s))] \right) - c \right\}.$$

By Lemma 2.3, the price $p_\iota(v, A)$ paid by the investor ι consists of a term entirely pinned down by the algorithm x , plus a term that does not depend on ι 's own report. Thus, for any two truthful mechanisms that use the same algorithm, (x, p) and (x, p') , the investor ι has the same privately optimal investments— $\text{BR}(x, p, \bar{\omega}) = \text{BR}(x, p', \bar{\omega})$ —so we henceforth suppress the payment rule argument p from $\text{BR}(\cdot)$.¹¹

The **welfare** of algorithm x at investment instance $\bar{\omega} = (g, I, \nu_{-\iota}, \mathcal{A})$ is

$$\bar{W}_x(\bar{\omega}) \equiv \min_{(\nu_\iota, c) \in \text{BR}(x, \bar{\omega})} \left\{ \left(\sum_{s \in S} g(s) W_x(\nu_\iota(s), \nu_{-\iota}(s), \mathcal{A}(s)) \right) - c \right\}.$$

¹⁰Negative costs $c < 0$ represent disinvestments compared to the status quo.

¹¹By assuming that the investor best-responds, we are abstracting from any computational limitations that the investor might face when there are many states.

We benchmark performance relative to the net welfare delivered by *ex ante* efficient investment and *ex post* efficient allocations. That is, the **optimal welfare** at investment instance $\bar{\omega} = (g, I, \nu_{-l}, A)$ is

$$\bar{W}^*(\bar{\omega}) \equiv \max_{(\nu_l, c) \in I} \left\{ \left(\sum_{s \in S} g(s) W^*(\nu_l(s), \nu_{-l}(s), \mathcal{A}(s)) \right) - c \right\}.$$

The benchmark $\bar{W}^*(\bar{\omega})$ is equal to the net welfare from selfish investment under a VCG mechanism. Given some $\beta \in [0, 1]$, algorithm x is a **β -approximation for investment** if for every investment instance $\bar{\omega}$, we have that $\bar{W}_x(\bar{\omega}) \geq \beta \bar{W}^*(\bar{\omega})$. Notably, since we are quantifying over $\iota \in N$ and I , this requires the inequality to hold regardless of which bidder is the investor and which investments are available.¹² We refer to the largest such β as the algorithm's **investment guarantee**.

Adding investment opportunities cannot improve an algorithm's performance.

Proposition 2.5. *If x is a β -approximation for investment, then x is a β -approximation for allocation.*

Proof. Any instance of the allocation problem (v, A) is equivalent to the instance of the investment problem $(g, I, \nu_{-l}, \mathcal{A})$ with the singleton investment technology $I = \{(\nu_l, 0)\}$, $\nu_l \equiv v_l$, $\nu_{-l} \equiv v_{-l}$, and $\mathcal{A} \equiv A$; the result then follows. \square

The converse of Proposition 2.5 does not hold in general—investment opportunities may strictly reduce the algorithm's performance guarantee. But when is an algorithm's allocative guarantee equal to its investment guarantee? We now determine the answer.

2.4.1 Reduction to the case without uncertainty

First, we simplify the problem by observing that, for our purposes, it is without loss of generality to focus on the case without uncertainty. That is, a **certain investment instance** $\bar{\omega}$ is an investment instance with just one state, so $|S| = 1$, and we abuse notation by writing such an instance as (I, v_{-l}, A) . An algorithm is a **β -approximation for certain investment** if $\bar{W}_x(\bar{\omega}) \geq \beta \bar{W}^*(\bar{\omega})$ for any certain investment instance $\bar{\omega}$.

The next theorem states that an algorithm's investment guarantee is unaffected by uncertainty.

¹²Our results extend naturally to the case in which some bidders are known in advance to be unable to make investments; in that case, our necessary conditions weaken to pertain only to those bidders who can make investments.

Theorem 2.6. *For any weakly monotone algorithm x and any $\beta \in [0, 1]$, x is a β -approximation for investment if and only if x is a β -approximation for certain investment.*

The intuition for Theorem 2.6 is as follows: Suppose we start from some investment instance with uncertainty. We can construct a related generalized investment instance where each investment is replaced by a generalized investment with the same values in every state but a *state-dependent* cost that makes the realized profit in each state equal to the original *ex ante* expected profit. This leaves the expected profits and expected costs from investing unchanged, so selfish investment in the original instance yields the same expected welfare as selfish investment in the new instance. If the algorithm x is a β -approximation for certain investment, then in the new instance, in every state, selfish investment achieves at least a fraction β of the welfare from the *ex post* efficient investment and the *ex post* efficient allocation. In turn, this is an upper bound for the expected welfare from the *ex ante* efficient investment and the *ex post* efficient allocation in the new instance, which is the same as in the original instance. It follows that x is a β -approximation for investment.

2.4.2 Performance under certain investment

Having reduced the problem with uncertain investment to the problem with certain investment, we now derive a necessary and sufficient condition for an algorithm x to be a β -approximation for certain investment.

We show a link between investment guarantees and algorithmic externalities. We can simplify the problem by focusing on the externalities that result from value changes in particular directions. Changing from v_n to \tilde{v}_n **confirms** outcome \tilde{o} if

$$[\tilde{v}_n - v_n] \cdot [\tilde{o} - o] \geq 0 \text{ for all outcomes } o. \quad (3)$$

Intuitively, (3) means that changing n 's value from v_n to \tilde{v}_n raises n 's value for \tilde{o} at least as much as it raises n 's value for any other outcome—equivalently, n 's marginal gain from switching from o to \tilde{o} does not fall. The system of inequalities (3) defines a convex cone with vertex at v_n . If x is weakly monotone, then any change from v_n to \tilde{v}_n that confirms $x_n(v, A)$ implies that

$$[\tilde{v}_n - v_n] \cdot [x_n(\tilde{v}_n, v_{-n}, A) - x_n(v, A)] = 0; \quad (4)$$

this follows by combining (1) and (3), with $\tilde{o} = x_n(v, A)$ and $o = x_n(\tilde{v}_n, v_{-n}, A)$. For any truthful (x, p) , type v_n cannot profitably imitate \tilde{v}_n and *vice versa*, so

$$v_n \cdot [x_n(\tilde{v}_n, v_{-n}, A) - x_n(v, A)] \leq p_n(\tilde{v}_n, v_{-n}, A) - p_n(v, A) \leq \tilde{v}_n \cdot [x_n(\tilde{v}_n, v_{-n}, A) - x_n(v, A)]. \quad (5)$$

From (4) and (5), it follows that

$$p_n(\tilde{v}_n, v_{-n}, A) - p_n(v, A) = v_n \cdot [x_n(\tilde{v}_n, v_{-n}, A) - x_n(v, A)]; \quad (6)$$

that is, the bidder with value v_n is indifferent between reporting v_n and reporting the confirming change \tilde{v}_n when facing (v_{-n}, A) .

The externalities from confirming changes reduce to a simple expression. In particular, they are equal to the difference between the welfare yielded by the new allocation at the old values and the welfare yielded by the old allocation at the old values.

Proposition 2.7. *For any weakly monotone x , any instance (v, A) , and any change from v_n to \tilde{v}_n that confirms $x_n(v, A)$, we have*

$$\mathcal{E}_x(\tilde{v}_n, (v, A)) = \sum_m [v_m \cdot [x_m(\tilde{v}_n, v_{-n}, A) - x_m(v, A)]]. \quad (7)$$

Proof. Substituting (6) into (2) yields (7). □

The key condition for our characterization is a lower bound on the externalities resulting from confirming value changes.

Definition 2.8. For some $\beta \in [0, 1]$, algorithm x has **β -bounded confirming externalities** if given any instance (v, A) and any change from v_n to \tilde{v}_n that confirms $x_n(v, A)$, we have

$$\mathcal{E}_x(\tilde{v}_n, (v, A)) \geq \beta W^*(v, A) - W_x(v, A). \quad (8)$$

The inequality (8) requires the algorithmic externality of the confirming change to exceed the lower bound, which is the negative of the slack in the allocative guarantee at instance (v, A) . Definition 2.8 is a necessary condition for an algorithm to be a β -approximation for certain investment, as we now prove.

Theorem 2.9. *For any weakly monotone algorithm x and any $\beta \in [0, 1]$, if x is a β -approximation for certain investment, then x has β -bounded confirming externalities.*

Proof. We prove the contrapositive. Suppose that x does not have β -bounded confirming externalities. Take any allocation instance (v, A) , bidder n , and value change \tilde{v}_n such that (8) does not hold, so

$$\mathcal{E}_x(\tilde{v}_n, (v, A)) < \beta W^*(v, A) - W_x(v, A). \quad (9)$$

We construct an investment instance with n as the investor by choosing an investment cost c so that n is indifferent between v_n at cost 0 and \tilde{v}_n at cost c , and denote the resulting

investment instance by $\bar{\omega}$. By construction, it is a best response for n to choose \tilde{v}_n at cost c , so we have

$$\bar{W}_x(\bar{\omega}) \leq W_x(\tilde{v}_n, v_{-n}, A) - c. \quad (10)$$

Moreover, the welfare of the best allocation when n chooses v_n at cost 0 is no more than the optimal benchmark at $\bar{\omega}$, so we have

$$W^*(v, A) \leq \bar{W}^*(\bar{\omega}). \quad (11)$$

Combining the inequalities (10), (9), and (11) yields

$$\begin{aligned} \bar{W}_x(\bar{\omega}) &\leq W_x(\tilde{v}_n, v_{-n}, A) - c \\ &= W_x(v, A) + \underbrace{\tilde{v}_n \cdot x_n(\tilde{v}_n, v_{-n}, A) - v_n \cdot x_n(v, A) - c}_{=p_n(\tilde{v}_n, v_{-n}, A) - p_n(v, A) \text{ by construction of } c} + \sum_{m \neq n} v_m \cdot [x_m(\tilde{v}_n, v_{-n}, A) - x_m(v, A)] \\ &= W_x(v, A) + \mathcal{E}_x(\tilde{v}_n, (v, A)) \\ &< \beta W^*(v, A) \\ &\leq \beta \bar{W}^*(\bar{\omega}), \end{aligned}$$

showing that x is not a β -approximation for certain investment. \square

We now state the converse of Theorem 2.9, i.e., that Definition 2.8 is a sufficient condition.

Theorem 2.10. *For any weakly monotone algorithm x and any $\beta \in [0, 1]$, if x has β -bounded confirming externalities, then x is a β -approximation for certain investment.*

The intuition for Theorem 2.10 is as follows: Algorithm x being a β -approximation for investment means the welfare of x at value profile v must not be too low. We characterize the exact bound, which depends on the optimal welfare only at the bidder's threshold value. Therefore x is a β -approximation for investment as long as a change from the threshold value to v_n has externalities that are not too negative. Since for any value v_n , there is value v'_n arbitrarily close to the threshold value such that the change from v'_n to v_n confirms $x_n(v'_n, v_{-n}, A)$, having β -bounded confirming externalities is also sufficient.

We summarize the preceding results in a corollary.

Corollary 2.11. *For any weakly monotone algorithm x and any $\beta \in [0, 1]$, the following statements are equivalent:*

1. x is a β -approximation for investment.
2. x is a β -approximation for certain investment.

3. x has β -bounded confirming externalities.

2.4.3 A tractable sufficient condition

Corollary 2.11 characterizes the allocation algorithms that attain performance guarantee β under investment. However, to check expression (8) in the definition of β -bounded confirming externalities, we must assess the *optimal* welfare W^* at some instance. Given that our main interest is in problems for which optimal allocations are hard to compute, verification of that condition may be intractable, so next we introduce the following more tractable sufficient condition.

Definition 2.12. Algorithm x **excludes confirming negative externalities** (“is **XCONE**”) if given any instance (v, A) and any change from v_n to \tilde{v}_n that confirms $x_n(v, A)$, we have

$$\mathcal{E}_x(\tilde{v}_n, (v, A)) \geq 0.$$

Theorem 2.13. For any weakly monotone algorithm x and any $\beta \in [0, 1]$, if x is **XCONE** and a β -approximation for allocation, then x is a β -approximation for investment.

Proof. Take any (v, A) and any change from v_n to \tilde{v}_n that confirms $x_n(v, A)$. We have

$$\mathcal{E}_x(\tilde{v}_n, (v, A)) \geq 0 \geq \beta W^*(v, A) - W_x(v, A);$$

the first inequality follows because x is **XCONE** and the second inequality follows because x is a β -approximation for allocation. Thus, we see that x has β -bounded confirming externalities. By Corollary 2.11, x is a β -approximation for investment. \square

Note that our results do not apply to non-deterministic algorithms, which randomize over multiple feasible allocations and have guarantees that hold only in expectation, nor to algorithms that apply to restricted value profiles, as our next example shows.

Example 2.14. Consider the randomized algorithm that packs a knapsack optimally or leaves it empty, each with probability $1/2$. This algorithm is a $1/2$ -approximation for allocation and has no externalities, so it is **XCONE**. Suppose there is just one bidder who can choose value 0 at cost 0 or value 3 at cost 2. In the threshold auction based on this approximation algorithm, the bidder will find that it is not profitable to invest, so the net welfare will be 0. But it is socially optimal for the bidder to invest and be packed, for net welfare 1. Thus, the randomized algorithm’s investment guarantee is 0. If we treat this as a deterministic algorithm that can choose between packing an item or half the item with half

the value, the same analysis shows that our result does not apply to problems with restricted value domains.

2.4.4 Relating XCONE to non-bossiness

Our XCONE condition is related to the standard mechanism design concept of non-bossiness. Algorithm x is **non-bossy** if $x_n(\tilde{v}_n, v_{-n}, A) = x_n(v, A)$ implies that $x(\tilde{v}_n, v_{-n}, A) = x(v, A)$, that is, if changing n 's value does not change n 's outcome, then it must not change others' outcomes, either. Algorithm x is **consistent** if (4) implies that $x_n(\tilde{v}_n, v_{-n}, A) = x_n(v, A)$; this holds, for instance, if whenever bidder n is indifferent between several outcomes at the threshold prices, the algorithm breaks ties according to some fixed order on outcomes.

Proposition 2.15. *In packing problems, every algorithm is consistent.*

Proof. We prove the contrapositive. Suppose $x_n(\tilde{v}_n, v_{-n}, A) \neq x_n(v, A)$; in a packing problem, this implies that $\tilde{v}_{n,\text{packed}} \neq v_{n,\text{packed}}$. Without loss of generality, suppose n is packed under $x_n(\tilde{v}_n, v_{-n}, A)$ and not packed under $x_n(v, A)$. Then the expression in the left-hand side of (4) is equal to $\tilde{v}_{n,\text{packed}} - v_{n,\text{packed}}$, which is non-zero (and hence (4) does not hold). \square

Proposition 2.16. *If x is weakly monotone, consistent, and non-bossy, then x is XCONE.*

Proof. Because x is weakly monotone, any change from v_n to \tilde{v}_n that confirms $x_n(v, A)$ implies (4). Next, because x consistent and non-bossy, (4) implies that $x(\tilde{v}_n, v_{-n}, A) = x(v, A)$. Thus, we have

$$0 = \sum_m v_m \cdot [x_m(\tilde{v}_n, v_{-n}, A) - x_m(v, A)] = \mathcal{E}_x(\tilde{v}_n, (v, A)),$$

where the final equality follows from Proposition 2.7. \square

2.4.5 Combinations of XCONE algorithms

A standard technique for addressing computationally difficult allocation problems is to run several candidate algorithms and select the best of their solutions; this can yield a better allocative guarantee than for each algorithm individually. However, the resulting algorithm may be bossy, even if the candidate algorithms are non-bossy.¹³ By contrast, if the candidate

¹³To see this, consider an allocation problem with three bidders; bidder n 's value for being packed is $v_{n,\text{packed}}$ and its value for not being packed is $v_{n,\text{unpacked}} = 0$. Algorithm x packs bidders 1 and 2 if $v_2^{\text{packed}} + 3 > v_3^{\text{packed}}$ and packs bidders 1 and 3 otherwise. Algorithm x' always packs just bidder 3. Let x'' select the best solution from x and x' . When $v_1^{\text{packed}} = 1$, $v_2^{\text{packed}} = 2$, $v_3^{\text{packed}} = 4$, x'' packs bidder 3, but if we raise v_3^{packed} to 8, then x'' packs bidders 1 and 3. Thus, while x and x' are non-bossy, x'' is bossy—yet all three algorithms are XCONE.

algorithms are XCONE, then the resulting algorithm is XCONE.

Proposition 2.17. *Let X be a collection of weakly monotone XCONE algorithms. If y is a weakly monotone algorithm that at each instance $(v, A) \in \Omega$ outputs a welfare-maximizing allocation from the collection $\{x(v, A)\}_{x \in X}$, then y is XCONE.¹⁴*

Proof. Consider any instance (v, A) and any \tilde{v}_n that confirms $y(v, A)$. Let $x \in X$ be such that $y(v, A) = x(v, A)$. Because x is weakly monotone and XCONE, Proposition 2.7 implies that

$$0 \leq \mathcal{E}_x(\tilde{v}_n, (v, A)) = \sum_m v_m \cdot [x_m(\tilde{v}_n, v_{-n}, A) - x_m(v, A)]; \quad (12)$$

hence, we have

$$\begin{aligned} \sum_m v_m \cdot y_m(v, A) &= \sum_m v_m \cdot x_m(v, A) \\ &\leq \sum_m v_m \cdot x_m(\tilde{v}_n, v_{-n}, A) \\ &\leq \sum_m v_m \cdot y_m(\tilde{v}_n, v_{-n}, A), \end{aligned} \quad (13)$$

where the first inequality follows from (12), and the second uses the definition of y . Rearranging (13) yields

$$\begin{aligned} 0 &\leq \sum_m v_m \cdot [y_m(\tilde{v}_n, v_{-n}, A) - y_m(v, A)] \\ &= \mathcal{E}_y(\tilde{v}_n, (v, A)), \end{aligned}$$

where the equality follows from Proposition 2.7 because y is weakly monotone. \square

Proposition 2.17 assumes that y is weakly monotone. Yet weak monotonicity of every algorithm in X does not necessarily imply weak monotonicity of y , even though y is a welfare-maximizing selection from X (see Example 2.19 below). One other advantage of XCONE algorithms is that such a y *does* inherit weak monotonicity from X when there are only two outcomes.

Proposition 2.18. *Suppose that $|O| = 2$, and let X be a collection of weakly monotone XCONE algorithms. If y is an algorithm that at each instance $(v, A) \in \Omega$ outputs a welfare-maximizing allocation from the collection $\{x(v, A)\}_{x \in X}$, then y is weakly monotone.*

¹⁴Our necessary and sufficient condition, Definition 2.8, also has this property. That is, if we replace the supposition that every algorithm in X is XCONE with the supposition that every algorithm in X has β -bounded confirming externalities, then a parallel proof yields the conclusion that y has β -bounded confirming externalities.

Proposition 2.18 does not generalize to $|O| > 2$. Indeed, there exist pairs of candidate algorithms, both weakly monotone and XCONE, such that the resulting y is not weakly monotone—as the following example illustrates.

Example 2.19. Consider an allocation problem with two bidders and three outcomes, and suppose that $V_1 = [0, 4] \times [0, 4] \times \{0\}$ and $V_2 = \{(5, 0, 0)\}$. We suppose that algorithm x always allocates outcome 2 to bidder 1 and outcome 3 to bidder 2, while algorithm \tilde{x} allocates outcome 1 to both bidders if $v_1 \geq 1$ and allocates outcome 3 to both bidders otherwise. Both x and \tilde{x} are weakly monotone and XCONE. Let y be an algorithm that outputs a welfare-maximizing allocation from the set $\{x(v_1, v_2), \tilde{x}(v_1, v_2)\}$. Under algorithm y , bidder 1 gets outcome 2 when $v_1 = (0, 1, 0)$ and outcome 1 when $v_1 = (2, 4, 0)$, so y is not weakly monotone.

2.4.6 Allowing multiple investors

Suppose that each bidder n has a finite set of feasible investments I_n and, as before, an investment consists of a function $\nu_n : S \rightarrow V_n$ and a cost $c \in \mathbb{R}$. Suppose that all bidders simultaneously choose investments, knowing that in each state $s \in S$, the resulting allocation and payments will be $x(\nu(s), A(s))$ and $p(\nu(s), A(s))$, for truthful mechanism (x, p) . The resulting investment game has a Nash equilibrium, possibly in mixed strategies.

Even for VCG mechanisms, not every Nash equilibrium of the investment game is efficient. Complementarities between the bidders can result in inefficient Nash equilibria, as the following example illustrates.

Example 2.20. Consider a packing problem with three bidders. It is feasible to pack any single bidder, or to pack bidder 2 and bidder 3 simultaneously. There is only one state and so no uncertainty: $|S| = 1$. Bidder 1 has a *status quo* value 10 for being packed, that is, its technology is the singleton $I_1 = \{(10, 0)\}$. Bidders 2 and 3 have the technology $I_2 = I_3 = \{(0, 0), (9, 1)\}$. Total welfare is maximized if both bidders 2 and 3 choose the investment $(9, 1)$, which leads to both being packed. However, if only one of them invests, then it is optimal to pack just Bidder 1. Under the VCG auction, there are two pure strategy Nash equilibrium investment profiles. In one Nash equilibrium, no bidder invests and Bidder 1 is packed, for net welfare 10. In the efficient Nash equilibrium, both Bidders 2 and 3 invest and both are packed, for net welfare 16.

Nevertheless, VCG mechanisms satisfy a different efficiency criterion: Conditional on any belief about the strategies of the other bidders, every best response for bidder n maximizes interim social welfare net of bidder n 's investment costs.

Our results extend this observation to include approximate efficiency. Any best response of bidder n to its belief about the other bidders' investments yields social welfare (net of n 's investment costs) that is at least a fraction β of what would be achieved by the interim efficient investment for bidder n and the *ex post* efficient allocation.

Proposition 2.21. *Let $h \in \Delta(I_{-n})$, with $h(\nu_{-n})$ denoting the marginal distribution. Let $(\nu_n, c) \in I_n$ be a best response for bidder n to the belief h given algorithm x . If x has β -bounded confirming externalities, then*

$$\left(\sum_{\nu_{-n}} h(\nu_{-n}) \sum_{s \in S} g(s) W_x(\nu_n(s), \nu_{-n}(s), \mathcal{A}(s)) \right) - c \geq \beta \max_{(\nu'_n, c') \in I_n} \left\{ \left(\sum_{\nu_{-n}} h(\nu_{-n}) \sum_{s \in S} g(s) W^*(\nu'_n(s), \nu_{-n}(s), \mathcal{A}(s)) \right) - c' \right\}.$$

Proof. Let us define a new single-investor instance that is payoff-equivalent for n , incorporating n 's belief h using an expanded state space $S \times S'$ and functions $\hat{\nu}_{-n} : S \times S' \rightarrow V_{-n}$. For each of bidder n 's investments $(\nu_n, c) \in I_n$, we define a corresponding investment $(\hat{\nu}_n, c)$ with $\hat{\nu}_n(s, s') \equiv \nu_n(s)$, and similarly define $\hat{\mathcal{A}}(s, s') \equiv \mathcal{A}(s)$. By Corollary 2.11, if x has β -bounded confirming externalities, then the desired inequality follows. \square

3 Application: Knapsack algorithms

The **knapsack problem** is a special case of the packing problem, in which each bidder n has possible values $V_{n,\text{packed}} = [0, \infty)$ and $V_{n,\text{unpacked}} = \{0\}$, each bidder has **size** $q_n \geq 0$, and the “knapsack” has **capacity** Q .

For knapsack problems, we abuse notation and use v_n to denote $v_{n,\text{packed}}$, bidder n 's value for being packed, since $v_{n,\text{unpacked}} \equiv 0$ uniformly. Without loss of generality, we also suppose that no bidder's size is more than Q .¹⁵ The set of feasible allocations is any subset of bidders $K \subseteq N$ such that $\sum_{n \in K} q_n \leq Q$. As before, let A denote the set of feasible allocations and let a be an element of A .

The knapsack problem is NP-Hard (Karp, 1972); there is no known polynomial-time algorithm that outputs optimal allocations (Cook, 2006; Fortnow, 2009).

3.1 Greedy algorithms

Dantzig (1957) suggested applying a **greedy algorithm** to the knapsack problem. For-

¹⁵Bidders with $q_n > Q$ can be deleted with no substantial change in an algorithm's runtime.

mally:

Algorithm (GREEDY). Sort bidders by the ratio of their values to their sizes so that

$$\frac{v_1}{q_1} \geq \frac{v_2}{q_2} \dots \geq \frac{v_{|N|}}{q_{|N|}}. \quad (14)$$

Add bidders to the knapsack one by one in the sorted order, so long as the sum of the sizes does not exceed the knapsack’s capacity. When encountering the first bidder that would violate the capacity constraint, stop.

Although the GREEDY algorithm performs well on some instances—including ones for which all bidders are small in relation to the capacity of the knapsack—its allocative guarantee is 0, as illustrated by the following example.

Example 3.1. Consider a knapsack with capacity 1 and two bidders. For some arbitrarily small $\varepsilon > 0$, let $v_1 = \varepsilon$, $q_1 = \frac{\varepsilon}{2}$, $v_2 = 1$, and $q_2 = 1$. The GREEDY algorithm picks bidder 1 and stops, whereas the optimal algorithm picks bidder 2. Thus, the worst-case performance of GREEDY is no better than ε of the optimum.

There is a standard modification of the GREEDY algorithm that improves the allocative guarantee for the knapsack problem (Williamson and Shmoys, 2011, p. 77). Let us define the “**smart greedy**” algorithm as follows.

Algorithm (SMARTGREEDY). Run the GREEDY algorithm. Compare the GREEDY algorithm’s packing to the packing that just packs the most valuable individual bidder, and output whichever has higher welfare.

SMARTGREEDY’s allocative guarantee is much better than GREEDY’s.

Proposition 3.2. SMARTGREEDY is a $\frac{1}{2}$ -approximation for allocation in the Knapsack problem.

Proof. For any instance ω , order the bidders by value/size as in (14). If GREEDY packs all bidders, then trivially $W^*(\omega) = W_{\text{SmartGreedy}}(\omega)$. Otherwise, let k be the lowest index of a bidder not packed by GREEDY and let K be the index of a bidder with maximum value. Optimal welfare $W^*(\omega)$ is no more than the best solution to the linear program in which we can pack fractional bidders, which—given that we have sorted the bidders in descending

order of value-to-size—in turn is no more than $\sum_{n=1}^k v_n$. It follows that

$$\begin{aligned}
W^*(\omega) &\leq \sum_{n=1}^k v_n \\
&= W_{\text{Greedy}}(\omega) + v_k \\
&\leq W_{\text{Greedy}}(\omega) + v_K \\
&\leq 2 \max \{W_{\text{Greedy}}(\omega), v_K\} \\
&= 2W_{\text{SmartGreedy}}(\omega);
\end{aligned}$$

hence, we see that SMARTGREEDY is a $\frac{1}{2}$ -approximation for allocation, as desired. \square

SMARTGREEDY is bossy, as our next example shows.

Example 3.3. Consider the knapsack instance with capacity 10 and three bidders; $v_1 = 2$, $v_2 = 1$, $v_3 = 8$, $q_1 = q_2 = 1$, and $q_3 = 9$. At this instance, SMARTGREEDY packs just bidder 3. If bidder 3 instead reports $v_3 = 10$, then SMARTGREEDY instead packs bidder 1 and bidder 3. Thus, we see that SMARTGREEDY is bossy. However, the adjustment just described is a confirming *positive* externality; raising the value of a packed bidder has strictly increased the welfare of other bidders.

The GREEDY and SMARTGREEDY algorithms are XCONE.

Proposition 3.4. *For the knapsack problem, the GREEDY algorithm and the SMARTGREEDY algorithm are both XCONE.*

Proof. Consider the bidders sorted by the GREEDY algorithm as in (14), and suppose the GREEDY algorithm packs bidders 1 through k . If we raise the value of a packed bidder (without changing sizes), then the GREEDY algorithm again packs bidders 1 through k . If we lower the value of an unpacked bidder, then the GREEDY algorithm terminates no earlier than before, packing at least bidders 1 through k . The only confirming externalities are positive ones; hence the GREEDY algorithm is XCONE.

Meanwhile, the algorithm that selects the most valuable single bidder is monotone and non-bossy and so is XCONE by Proposition 2.16, as well. Thus, by Proposition 2.18, the SMARTGREEDY algorithm is monotone, and so by Proposition 2.17 it is XCONE. \square

Proposition 3.2, Proposition 3.4, and Theorem 2.13 yield the following corollary.

Corollary 3.5. *The SMARTGREEDY algorithm is a $\frac{1}{2}$ -approximation for investment.*

The SMARTGREEDY algorithm has both confirming externalities (Example 3.3) and negative externalities, as the next example demonstrates. Crucially, it has no confirming negative externalities.

Example 3.6. Suppose we have three bidders with sizes $(.5, .5, .6)$ and values $(1, 1, 0)$, and a knapsack with capacity 1. The SMARTGREEDY algorithm packs the first two bidders. Raising the third bidder’s value to 2 raises its payment by 1.2 but reduces the welfare of the other bidders by 2. However, this value change is not confirming.

3.2 Fully polynomial-time approximation schemes

State-of-the-art knapsack algorithms that run in polynomial time have stronger allocative guarantees than the SMARTGREEDY algorithm. Can fast XCONE algorithms be constructed that match their performance? Or does restricting attention to XCONE algorithms force us to accept slower speeds or poorer allocative guarantees? We will answer these questions shortly.

Our construction in the sequel can be read at two different levels. To follow in full detail, readers should be acquainted with the theory of computation—in particular with how instances are represented as input strings and how running time is defined as a function of input size.¹⁶ Alternatively, readers can follow the proofs that our new algorithms are XCONE while observing that they inherit their allocative guarantees and polynomial runtimes from the other algorithms used in the construction.

As is standard for computational running-time analyses, we now assume that the bidders’ values are non-negative integers.¹⁷ Under this assumption, the input size is polynomial in the logarithm of the highest value $\log(\max_n \{v_n\})$ and the number of bidders $|N|$.¹⁸

We have used the SMARTGREEDY algorithm for illustration, but there are fast knapsack algorithms that do better. In particular, there exist families of algorithms indexed by parameter $\varepsilon > 0$, that yield $(1 - \varepsilon)$ -approximations for allocation, with running time polynomial in ε^{-1} and the input size. Such a family is called a **fully polynomial-time approximation scheme** (FPTAS).

Briest et al. (2005) (henceforth BKV) constructed a weakly monotone FPTAS for the knapsack problem. Our construction modifies two steps in theirs to ensure that the algorithms have the XCONE property in addition to being a weakly monotone FPTAS.

¹⁶These formalisms can be reviewed in the Arora and Barak (2009, pp. 9–37) textbook.

¹⁷Real numbers can take infinitely many bits to represent, complicating statements about input size. But note that the restriction to non-negative integers is only needed for our running-time analysis—the algorithm works fine for non-integral values, which is crucial for satisfying our product structure assumption.

¹⁸It is conventional to take the logarithm with base 2, but this statement is true for any base.

Suppose we have some allocation instance with value profile v , and our desired allocative guarantee is $(1 - \varepsilon)$. The first step of the BKV construction is to round each value to a grid. We define a family of modified value profiles, one for each non-negative integer $\ell \in \mathbb{N}$, essentially, censoring values above $2^{\ell+1}$ and then rounding the values to a grid with step size $\gamma^{\varepsilon, \ell} := \frac{\varepsilon 2^\ell}{|N|}$. Formally, for given $\varepsilon > 0$ and ℓ , let us define a modified value profile $v^{\varepsilon, \ell}$ as follows:

1. $v'_n := \min\{v_n, 2^{\ell+1}\}$ (for all n).
2. $v_n^{\varepsilon, \ell} := \lfloor v'_n / \gamma^{\varepsilon, \ell} \rfloor \cdot \gamma^{\varepsilon, \ell}$ (for all n).

An exact optimum for the modified values $v^{\varepsilon, \ell}$ can be computed in polynomial time (for the textbook algorithm, see [Williamson and Shmoys \(2011, pp. 65–68\)](#)).

Let $x^*(\tilde{v})$ be any selection from the set of optimal allocations at value profile \tilde{v} (implicitly, given the feasible allocations A), that is $x^*(\tilde{v}) \in \operatorname{argmax}_{a \in A} \{\sum_n [\tilde{v}_n \cdot a_n]\}$. Given parameter $\varepsilon > 0$, the BKV allocation rule (henceforth **the BKV rule**) selects an allocation in the infinite set

$$\{x^*(v^{\varepsilon, \ell})\}_{\ell \in \mathbb{N}}$$

that maximizes performance according to the modified values, i.e., $\max_\ell \{\sum_n [v_n^{\varepsilon, \ell} \cdot x_n^*(v^{\varepsilon, \ell})]\}$. One needs to search only a finite number of non-negative integers ℓ to find the desired maximum, because for large enough ℓ all the modified values round to 0. Combining the preceding steps, we obtain an algorithm that computes the BKV rule in polynomial time.

Proposition 3.7 ([Briest et al. \(2005\)](#)). *The BKV rule is weakly monotone, a $(1 - \varepsilon)$ -approximation for allocation, and can be computed in $\operatorname{poly}(\varepsilon^{-1}, |N|, \log(\max_n \{v_n\}))$ time.*

Despite the appealing properties described in [Proposition 3.7](#), the BKV rule has confirming negative externalities because a large investment can make only large values of ℓ relevant in the preceding computation, reducing the total welfare of the other bidders. For sufficiently large investments, this negative confirming externality can be arbitrarily bad, as we now state formally.

Proposition 3.8. *For all $\delta > 0$, there exists $\varepsilon < \delta$ such that the BKV rule with parameter ε has an investment guarantee of 0.*

The proof of [Proposition 3.8](#) uses an example that mimics the satisficing example from the Introduction, but using the more complicated FPTAS algorithm. The two examples share these key properties:

1. When bidder 1 does not invest, its value is very low so that nearly the entire value is derived from the packing of other bidders.
2. When bidder 1 does invest, its value becomes very high but it incurs an equally high cost, so the investment is barely profitable. The algorithm then packs bidder 1 but obtains almost no value from the other bidders, so the value of the packing net of investment cost falls to zero.

Nevertheless, we can construct a XCONE FPTAS by modifying the BKV rule in two ways. First, instead of defining x^* to be an arbitrary maximizer when there are multiple maximizers, we limit the selection so that we never pack any bidders whose values are exactly 0, and we break ties among maximizers using a strict ordering. The resulting x^* is non-bossy. Second, where BKV selects an allocation in the family $\{x^*(v^{\varepsilon, \ell})\}_{\ell \in \mathbb{N}}$ that maximizes welfare using the *modified* values $\sum_n [v_n^{\varepsilon, \ell} \cdot a_n]$, our modification selects an allocation in the same family that maximizes welfare using the *actual* values $\sum_n [v_n \cdot a_n]$. For any parameter $\varepsilon > 0$, we use \check{x}^ε to denote the resulting allocation rule.

Proposition 3.9. *The allocation rule \check{x}^ε is XCONE, weakly monotone, and a $(1 - \varepsilon)$ -approximation for allocation.*

Proof. Let us define the allocation rule

$$\check{x}^{\varepsilon, \ell}(v) \equiv x^*(v^{\varepsilon, \ell}).$$

The allocation rule x^* as modified is weakly monotone, and the censoring and rounding operations are monotone transformations, so $\check{x}^{\varepsilon, \ell}$ is weakly monotone. Moreover, x^* is non-bossy, so $\check{x}^{\varepsilon, \ell}$ is non-bossy. Thus, $\check{x}^{\varepsilon, \ell}$ is XCONE by Proposition 2.15 and Proposition 2.16.

Now, \check{x}^ε chooses the best output from the collection $\{\check{x}^{\varepsilon, \ell}(v)\}_{\ell \in \mathbb{N}}$, so \check{x}^ε is weakly monotone by Proposition 2.18. Next, applying Proposition 2.17 yields the conclusion that \check{x}^ε is XCONE.

The BKV rule is a $(1 - \varepsilon)$ -approximation for allocation and chooses the allocation from the collection $\{\check{x}^{\varepsilon, \ell}(v)\}_{\ell \in \mathbb{N}}$ that maximizes welfare under the modified values. The allocation $\check{x}^\varepsilon(v)$ is selected from the same collection to maximize welfare under the actual values, so it achieves a weakly higher welfare than the BKV rule. Hence \check{x}^ε is a $(1 - \varepsilon)$ -approximation for allocation. \square

Corollary 3.10. *The allocation rule \check{x}^ε is a $(1 - \varepsilon)$ -approximation for investment.*

Moreover, since \check{x}^ε is computed by tweaking the BKV algorithm, it inherits BKV's polynomial time property, resulting in a FPTAS.

Proposition 3.11. *The allocation rule \check{x}^ε can be computed in $\text{poly}(\varepsilon^{-1}, |N|, \log(\max_n \{v_n\}))$ time.*

Proposition 3.9 and Proposition 3.11 demonstrate that good investment guarantees and efficient computations are sometimes compatible: there is a FPTAS that achieves both. Although the BKV FPTAS and our modification both run in polynomial time, ours requires additional steps: our FPTAS is slower than the BKV FPTAS. Further details are in the proof of Proposition 3.11.

We have focused on the knapsack problem for ease of exposition, but BKV showed how to construct a monotone FPTAS for a range of weakly NP-complete problems, such as job scheduling with deadlines and the constrained shortest-path problem.¹⁹ Our method adapts easily to convert the BKV FPTAS to a XCONE FPTAS for those problems as well.

4 Discussion

Mechanism design analysis in economics has traditionally focused on mechanisms that exactly optimize some objective like welfare, revenue, or consumer surplus, neglecting issues of computational hardness. Yet exact optimization is tractable only for small problems or problems with special structure.

Practical mechanisms without optimization can be created by using the large corpus of fast approximation algorithms developed by computer scientists, but doing so raises new questions. Approximation algorithms have heretofore been designed for short-run problems in which participants' values are fixed exogenously, but in practice, participants can often make *ex ante* investments that alter their values. In this paper, we study investment incentives in a class of environments in which there is a finite number of outcomes for each bidder and the bidder's possible values lie in a product of intervals. We asked three general questions:

1. *Can mechanisms based on (deterministic) approximation algorithms avoid distorting participants' investment incentives as VCG mechanisms, based on optimization, do?*
2. *When do such mechanisms have the same allocative and investment guarantees?*
3. *Is there a trade-off between an algorithm's speed, its allocative guarantee, and the requirement that its performance extends to the case when an agent can invest?*

¹⁹If an allocation problem is NP-complete, but one can find an optimal allocation with running time polynomial in the *numeric value* of the largest integer in the input, then it is called "weakly NP-complete." Note that such an algorithm might still run in time exponential in the *length* of the input.

To frame the first question, we began by showing that the externalities from any truthful mechanism depend only on the algorithm, and not on which prices are used to promote truthful reporting. For that reason, we call these “algorithmic externalities.” Then, for the first question, we find a negative answer: unless the algorithm mimics welfare maximization on some possibly limited set of allocations, there are necessarily non-zero externalities that can cause privately profitable investments to reduce welfare or welfare-increasing investments to be privately unprofitable.

Our analysis of investment guarantees hinges on a new category of externalities that we dub “confirming” algorithmic externalities. These arise when a change in a bidder’s report that raises the relative value of its outcome results in an externality to other bidders. We show that an algorithm’s worst-case allocative guarantee extends to become an investment guarantee if and only if its confirming negative externalities are not too large. That condition, however, can be hard to verify, so we also offer a sufficient condition—XCONE—that can be easier to check. An XCONE algorithm is one that excludes confirming negative externalities, but may have confirming positive externalities. We show that for some algorithms for the knapsack problem including GREEDY and SMARTGREEDY, the XCONE condition can be checked and verified without much difficulty. However, the XCONE condition also fails for some algorithms with very good—even arbitrarily good—performance for the short-run allocation problem.

Towards the third question, we study one particular FPTAS for the knapsack problem, modifying it with investment incentives in mind. The result is a new XCONE FPTAS—a collection of algorithms that, for every ε , is XCONE, always achieves at least a $1 - \varepsilon$ fraction of the optimum, and runs in time that is polynomial in the size of the input and ε^{-1} .

More broadly, there is a long tradition in economics of studying the performance of a competitive equilibrium, which assumes that all decisions, short-run and long-run, are guided by optimization. Because some problems are too complex for optimization, it can be valuable to study economies in which approximation algorithms replace optimization. As we have shown, investment incentives in economies based on approximation algorithms can differ sharply from economies with full optimization, and approximations can also affect how participants understand mechanisms in practice, create new opportunities for coordination or collusion, and influence post-auction resale markets. Given the close connection between weakly monotone algorithms and truthful mechanisms, it seems possible—and important—to analyze how these and other economic properties of mechanisms reflect properties of their underlying algorithms.

References

- ARORA, S. AND B. BARAK (2009): *Computational Complexity: A Modern Approach*, Cambridge University Press.
- AROZAMENA, L. AND E. CANTILLON (2004): “Investment incentives in procurement auctions,” *Review of Economic Studies*, 71, 1–18.
- BERGEMANN, D. AND J. VÄLIMÄKI (2002): “Information Acquisition and Efficient Mechanism Design,” *Econometrica*, 70, 1007–1033.
- BIENSTOCK, D. AND A. VERMA (2019): “Strong NP-hardness of AC power flows feasibility,” *Operations Research Letters*, 47, 494–501.
- BIKHCHANDANI, S., S. CHATTERJI, R. LAVI, A. MU’ALEM, N. NISAN, AND A. SEN (2006): “Weak Monotonicity Characterizes Deterministic Dominant-Strategy Implementation,” *Econometrica*, 74, 1109–1132.
- BIKHCHANDANI, S., S. DE VRIES, J. SCHUMMER, AND R. V. VOHRA (2011): “An ascending vickrey auction for selling bases of a matroid,” *Operations Research*, 59, 400–413.
- BRIEST, P., P. KRISTA, AND B. VÖCKING (2005): “Approximation techniques for utilitarian mechanism design,” in *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*, 39–48.
- COOK, S. (2006): “The P versus NP problem,” in *The Millennium Prize Problems*, ed. by J. A. Carlson, A. Jaffe, and A. Wiles, American Mathematical Society Providence, 87–104.
- DANTZIG, G. B. (1957): “Discrete-variable extremum problems,” *Operations Research*, 5, 266–288.
- DOBZINSKI, S. AND N. NISAN (2007): “Mechanisms for multi-unit auctions,” in *Proceedings of the 8th ACM Conference on Electronic Commerce*, 346–351.
- DUGHMI, S., J. D. HARTLINE, R. KLEINBERG, AND R. NIAZADEH (2017): “Bernoulli factories and black-box reductions in mechanism design,” in *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, 158–169.
- FORTNOW, L. (2009): “The status of the P versus NP problem,” *Communications of the ACM*, 52, 78–86.

- GERSHKOV, A., B. MOLDOVANU, P. STRACK, AND M. ZHANG (2021): “A Theory of Auctions with Endogenous Valuations,” *Journal of Political Economy*, 129, 1011–1051.
- HARTLINE, J. D. (2016): *Mechanism Design and Approximation*, mimeo.
- HARTLINE, J. D. AND B. LUCIER (2015): “Non-optimal mechanism design,” *American Economic Review*, 105, 3102–24.
- HATFIELD, J. W., F. KOJIMA, AND S. D. KOMINERS (2014): “Investment Incentives in Labor Market Matching,” *American Economic Review Papers & Proceedings*, 104, 436–441.
- (2019): “Strategy-Proofness, Investment Efficiency, and Marginal Returns: An Equivalence,” Becker Friedman Institute Working Paper.
- HOLZMAN, R., N. KFIR-DAHAV, D. MONDERER, AND M. TENNENHOLTZ (2004): “Bundling equilibrium in combinatorial auctions,” *Games and Economic Behavior*, 47, 104–123.
- KARP, R. M. (1972): “Reducibility among combinatorial problems,” in *Complexity of Computer Computations*, ed. by R. E. Miller and J. W. Thatcher, Springer, 85–103.
- LAVAEI, J. AND S. H. LOW (2011): “Zero duality gap in optimal power flow problem,” *IEEE Transactions on Power Systems*, 27, 92–107.
- LAVI, R., A. MU’ALEM, AND N. NISAN (2003): “Towards a characterization of truthful combinatorial auctions,” in *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, 574–583.
- LEHMANN, D., R. MÜLLER, AND T. SANDHOLM (2006): “The Winner Determination Problem,” in *Combinatorial Auctions*, ed. by P. Cramton, Y. Shoham, and R. Steinberg, MIT Press, 297–318.
- LEHMANN, D., L. I. O’CALLAGHAN, AND Y. SHOHAM (2002): “Truth revelation in approximately efficient combinatorial auctions,” *Journal of the ACM*, 49, 577–602.
- LEYTON-BROWN, K., P. MILGROM, AND I. SEGAL (2017): “Economics and computer science of a radio spectrum reallocation,” *Proceedings of the National Academy of Sciences*, 114, 7202–7209.
- MILGROM, P. (2017): *Discovering Prices*, Columbia University Press.

- MILGROM, P. AND I. SEGAL (2002): “Envelope theorems for arbitrary choice sets,” *Econometrica*, 70, 583–601.
- (2020): “Clock auctions and radio spectrum reallocation,” *Journal of Political Economy*, 128, 1–31.
- MU’ALEM, A. AND N. NISAN (2008): “Truthful approximation mechanisms for restricted combinatorial auctions,” *Games and Economic Behavior*, 64, 612–631.
- NISAN, N. AND A. RONEN (1999): “Algorithmic mechanism design,” in *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing*, 129–140.
- (2007): “Computationally feasible VCG mechanisms,” *Journal of Artificial Intelligence Research*, 29, 19–47.
- PARDALOS, P. M., D.-Z. DU, AND R. L. GRAHAM (2013): *Handbook of Combinatorial Optimization*, Springer.
- ROGERSON, W. P. (1992): “Contractual solutions to the hold-up problem,” *Review of Economic Studies*, 59, 777–793.
- SAKS, M. AND L. YU (2005): “Weak monotonicity suffices for truthfulness on convex domains,” in *Proceedings of the 6th ACM Conference on Electronic Commerce*, 286–293.
- SANDHOLM, T. (2002): “Algorithm for optimal winner determination in combinatorial auctions,” *Artificial Intelligence*, 135, 1–54.
- TOMOEDA, K. (2019): “Efficient investments in the implementation problem,” *Journal of Economic Theory*, 182, 247–278.
- VICKREY, W. (1961): “Counterspeculation, auctions, and competitive sealed tenders,” *The Journal of Finance*, 16, 8–37.
- WILLIAMSON, D. P. AND D. B. SHMOYS (2011): *The Design of Approximation Algorithms*, Cambridge University Press.

A Proofs omitted from the main text

Preliminaries and notation

Before getting into the proofs, we introduce a few notations and conventions. The value profile for bidders other than n , v_{-n} ; the set of possible allocations, A ; the algorithm, x ; and the probability distribution, g , usually do not change within a given proof. Therefore, we often suppress the dependence on these parameters to ease notation (see Table 2).

Meanwhile, truthfulness of the mechanism (x, p) implies that for every allocation instance (v, A) , x assigns each bidder n an outcome that maximizes its value minus its threshold price. We call this maximum the bidder's **normalized utility** and denote it as follows:

$$u_n(v, A, x) \equiv [v_n - \tau_n(v_{-n}, A, x)] \cdot x_n(v, A) = \max_{o \in O} \{v_{n,o} - \tau_{n,o}(v_{-n}, A, x)\}.$$

The normalized utility corresponds to the bidder's utility in the mechanism with prices $p_n(v, A) = \tau_n(v_{-n}, A, x) \cdot x_n(v, A)$; other truthful mechanisms may shift prices and utility by a strategically irrelevant additive term. By construction, we have $u_n(v, A, x) \geq 0$.

We extend the normalized utility notation u to the case of investment as follows:

$$u_\iota((v_\iota, c), v_{-\iota}, A, x) \equiv \max_{o \in O} \{v_{\iota,o} - \tau_{\iota,o}(v_{-\iota}, A, x)\} - c,$$

$$u_\iota(I, v_{-\iota}, A, x) \equiv \max_{(v_\iota, c) \in I} \{u_\iota((v_\iota, c), v_{-\iota}, A, x)\}.$$

This allows us to talk about normalized utility for an investor facing a cost or set of investment opportunities.

Notation	Description	Suppressed Notation
$x(v_n, v_{-n}, A)$	allocation of algorithm x	$x(v_n)$
$W_x(v_n, v_{-n}, A)$	welfare of algorithm x	$W_x(v_n)$
$W^*(v_n, v_{-n}, A)$	welfare of the optimal algorithm	$W^*(v_n)$
$\text{BR}(x, g, I, v_{-n}, A)$	best response for the investor	$\text{BR}(x, I)$
$\tau_n(v_{-n}, A, x)$	threshold price	τ_n
$u_n(v_n, v_{-n}, A, x)$	normalized utility of bidder n	$u_n(v_n)$
$u_\iota((v_\iota, c), v_{-\iota}, A, x)$	normalized utility of bidder ι	$u_\iota(v_\iota, c)$
$u_\iota(I, v_{-\iota}, A, x)$	normalized utility of bidder ι	$u_\iota(I)$
$p_n(v_n, v_{-n}, A)$	price paid by bidder n	$p_n(v_n)$
$\mathcal{E}_x(\tilde{v}_n, (v, A))$	externality from changing v_n to \tilde{v}_n	$\mathcal{E}_x(\tilde{v}_n, v_n)$

Table 2: Correspondence table for the simplified notation we often use in this section

Proof of Theorem 2.4

Theorem 2.4. *Suppose that algorithm x is weakly monotone. Then x has no algorithmic externalities ($\mathcal{E}_x \equiv 0$) if and only if x is welfare-equivalent to a maximal-in-range algorithm.*

If x is welfare-equivalent to a maximal-in-range algorithm with some range R , then the welfare coincides with that of the VCG mechanism restricted to R . Since a VCG mechanism has zero externalities, x also has zero externalities because of the following lemma.

Lemma A.1. *If algorithms x and x' are weakly monotone and welfare-equivalent, then they have the same threshold prices and externalities.*

Proof. Assume x and x' are welfare-equivalent and have different threshold prices, say $\tau_{n,o}(x) < \tau_{n,o}(x')$. Then there exists some value v_n with $v_{n,o} < \tau_{n,o}(x')$ and $x_n(v_n) = o$. Let us define v'_n with $v'_{n,o} \in [v_{n,o}, \tau_{n,o}(x'))$, and $v'_{n,o'} = v_{n,o'}$ for $o' \neq o$. By (weak) monotonicity, $x_n(v'_n) = o$ —but by the definition of the threshold price, $x'_n(v'_n) \neq o$. However, while there are infinitely many such $v'_{n,o} \in [v_{n,o}, \tau_{n,o}(x'))$, there are only a finite set of values $v'_{n,o}$ where an allocation a with $a_n = o$ can have the same welfare as an allocation a' with $a'_n \neq o$. Thus, we obtain a contradiction to the hypothesis that x and x' are welfare-equivalent. Therefore x and x' must have the same threshold prices. Now,

$$\begin{aligned} \mathcal{E}_{x,p}(\tilde{v}_n, v_n) &\equiv \underbrace{p_n(\tilde{v}_n) - p_n(v_n)}_{\text{change in } n\text{'s threshold payment}} + \underbrace{\sum_{m \neq n} [v_m \cdot [x_m(\tilde{v}_n) - x_m(v_n)]]}_{\text{effect on others' welfare}} \\ &= W_x(\tilde{v}_n) - u_n(\tilde{v}_n) - W_x(v_n) + u_n(v_n) \\ &= W_x(\tilde{v}_n) - \max_{o \in O} \{\tilde{v}_{n,o} - \tau_{n,o}\} - W_x(v_n) + \max_{o \in O} \{v_{n,o} - \tau_{n,o}\}, \end{aligned}$$

so the externalities of x only depend on the welfare and threshold price functions—and similarly for x' .²⁰ Thus, since x and x' welfare-equivalent (by hypothesis) and have the same threshold prices (by the first part of the lemma), we see that they have the same externalities. \square

Let $W(v, a) \equiv \sum_n [v_n \cdot a_n]$ denote the welfare of allocation a at a value profile v . We write $a \simeq a'$ if $W(v, a) = W(v, a')$ for every value profile v . Define the **modified domain** as $D = \{v \in V : W(v, a) = W(v, a') \implies a \simeq a'\}$ and the **modified range** as $R = x(D)$.

We now fix a weakly monotone algorithm x with no algorithmic externalities. In outline, our proof of the forward direction of Theorem 2.4 constructs a maximal-in-range algorithm

²⁰Here we use the normalized utility notation introduced in the Preliminaries & Notation section at the start of the appendix.

x' welfare-equivalent to x as follows. We have just defined a subdomain of values $D \subseteq V$ on which no two essentially different allocations have the same welfare. Any algorithm x' that is welfare-equivalent to x must satisfy $x'(v) = x(v)$ for $v \in D$. For $v \notin D$, we use the monotonicity and zero-externality properties of x to show that there exists $a \in R$ such that $W(v, a) = W(v, x(v))$ and set $x'(v) = a$. This ensures that x and x' are welfare-equivalent. We then use the same properties to establish that x' is a maximal-in-range algorithm with range R , which finishes the proof.

Claim A.2. *Algorithm x is welfare-equivalent to a maximal-in-range algorithm with range R .*

Proof. The proof relies on the following two lemmata. The first characterizes properties of the welfare function for an algorithm that has no externalities. The second shows that the modified domain D is dense.

Lemma A.3. *If algorithm x has no externalities, then W_x is*

- *non-decreasing in v_n and*
- *1-Lipschitz in v_n*

in the sup norm.

Proof. Recall that the formula for externalities is

$$\begin{aligned} \mathcal{E}_{x,p}(\tilde{v}_n, (v, A)) &\equiv \underbrace{p_n(\tilde{v}_n, v_{-n}, A) - p_n(v, A)}_{\text{change in } n\text{'s threshold payment}} + \underbrace{\sum_{m \neq n} [v_m \cdot [x_m(\tilde{v}_n, v_{-n}, A) - x_m(v, A)]]}_{\text{effect on others' welfare}} \\ &= W_x(\tilde{v}_n) - u_n(\tilde{v}_n) - W_x(v_n) + u_n(v_n). \end{aligned}$$

Since x has no externalities, $\mathcal{E}_{x,p}(\tilde{v}_n, (v, A)) = 0$, so $W_x(v_n) - u_n(v_n) = W_x(\tilde{v}_n) - u_n(\tilde{v}_n)$. Therefore, $W_x(\cdot)$ is a constant plus $u_n(\cdot)$. It is clear that $u_n(\cdot)$ is nondecreasing and 1-Lipschitz; hence we see that W_x is nondecreasing and 1-Lipschitz, as well. \square

Lemma A.4. *The modified domain D is non-empty. In addition, for all $v \in V$, $d \in D$, and $\varepsilon > 0$, there exists $v' \in D$ such that*

- $\|v' - v\| < \varepsilon$, and
- $\prod_{n,o} \{v'_{n,o}, d_{n,o}\} \subseteq D$.

Proof. To show D is non-empty, we construct a value profile $d \in D$. We describe its components $d_{n,o}$ by ordering the pairs (n, o) , beginning with the pairs (n, o) for which $V_{n,o}$ is a singleton and listing the rest in arbitrary order, indexed by k . In step $k = 0$, for each pair (n, o) such that $V_{n,o}$ is a singleton, we fix $d_{n,o}$ to be equal to the sole element of $V_{n,o}$. Set $B_0 := \{d_{n,o} \in V_{n,o} : V_{n,o} \text{ is a singleton}\}$ (and $B_0 = \emptyset$ if no $V_{n,o}$ is a singleton). For each step $k \geq 1$, given a finite set B_{k-1} and any two subsets $B', B'' \subseteq B_{k-1}$, consider these equations:

$$d_{n,o} + \sum_{b \in B'} b = \sum_{b \in B''} b.$$

There are finitely many such equations and, for $k \geq 1$, the interval $V_{n,o}$ has non-empty interior, so there exists some $d_{n,o} \in V_{n,o}$ that satisfies none of the equations (15). We set $B_k := B_{k-1} \cup \{d_{n,o}\}$ and iterate until d has been constructed. Suppose that allocations a and a' satisfy $W(d, a) = W(d, a')$. Then, for all n , either $a_n = a'_n$ or V_{n,a_n} and V_{n,a'_n} are both singletons, so $a \simeq a'$ and hence $d \in D$.

The second half of the lemma is proved by constructing v' in a similar way. Call a value v' **generic (with respect to d)** if $\prod_{n,o} \{v'_{n,o}, d_{n,o}\} \subseteq D$. Start with $v' = d$. If $V_{n,o}$ is a singleton, then $v'_{n,o} = v_{n,o}$. Otherwise, $V_{n,o}$ is an interval with non-empty interior, so there exists some value $v'_{n,o}$ within $\frac{\varepsilon}{|N| \cdot |O|}$ of $v_{n,o}$ that keeps v' generic (because only a finite number of choices for $v'_{n,o}$ would result in v' being non-generic). \square

First we show that at any value profile v , the welfare of x is less than or equal to the welfare of some allocation in the modified range R . Because the welfare $W(v, a)$ is continuous in v and the set D is dense in V , the welfare of x at any value profile (even outside of D) must be equal to the welfare of some allocation in R at that value profile.

Formally, for any $\varepsilon > 0$ and $v \in V$, by Lemma A.4, we can find $v' \in D$ so that $\|v' - v\| < \varepsilon$. By Lemma A.3, we get

$$|W_x(v) - W_x(v')| < \varepsilon$$

and we have

$$|W_x(v') - W(v, x(v'))| < \varepsilon.$$

By the triangle inequality, we then have

$$W_x(v) < W(v, x(v')) + 2\varepsilon \leq \left(\max_{r \in R} \{W(v, r)\} \right) + 2\varepsilon.$$

Thus, we have

$$W_x(v) \leq \max_{r \in R} \{W(v, r)\}. \tag{15}$$

Now we show that the welfare of x at a value profile v is bounded below by the welfare of any allocation $r \in R$ at that value profile. This is because if we start at a value profile d where the allocation is r and look at the welfare as we change the value profile to v , weak monotonicity ensures certain changes do not change the allocation r , and the welfare function being 1-Lipschitz ensures that the other changes weakly increase the algorithm's welfare compared to the welfare of allocation r .

Formally, for any $r \in R$, let $d \in D$ be a value profile where $r = x(d)$. For any $v \in V$ and $\varepsilon > 0$, we can construct v' as in Lemma A.4. Consider the profile v'' where

$$v''_{n,o} = \begin{cases} \max\{v'_{n,o}, d_{n,o}\} & x_n(d) = o \\ \min\{v'_{n,o}, d_{n,o}\} & x_n(d) \neq o. \end{cases}$$

We now prove that $x(v'') \simeq r$. Consider changing the value profile from d to v'' one element at a time, and let v^k denote the value profile in step k ; we show that at each step we have $x(v^k) \simeq r$. We argue by contradiction; suppose at some step we are at value profile v^k and $x(v^k) \simeq r$, but $x(v^{k+1}) \not\simeq r$. We define $\tilde{\alpha} \equiv \inf\{\alpha \in [0, 1] : x(\alpha v^{k+1} + (1 - \alpha)v^k) \not\simeq r\}$, and $\tilde{v} \equiv \tilde{\alpha}v^{k+1} + (1 - \tilde{\alpha})v^k$. We can find an allocation $a \not\simeq r$ and a sequence of value profiles $(v^l)_{l=1}^\infty$ such that $x(v^l) = a$, v^l is a convex combination of v^k and v^{k+1} , and $\lim_{l \rightarrow \infty} v^l = \tilde{v}$. By continuity of W_x (from Lemma A.3), we have

$$W(\tilde{v}, a) = \lim_{l \rightarrow \infty} W_x(v^l) = W_x(\tilde{v}) = \lim_{l \rightarrow \infty} W_x((1/l)v^k + (1 - 1/l)\tilde{v}) = W(\tilde{v}, r). \quad (16)$$

Moreover because x is weakly monotone, we have $a_{n,o} = r_{n,o}$ for the n, o pertaining to step k . The value profiles v^k and \tilde{v} are identical except for the element pertaining to n, o , so (16) implies that $W(v^k, a) = W(v^k, r)$. But since the value profile v^k is in D (by construction of v' and v''), it must be that $a \simeq r$, a contradiction. We have proven that $x(v'') \simeq r$, so

$$W_x(v'') = W(v'', r).$$

By Lemma A.3, the welfare function is nondecreasing and 1-Lipschitz, so²¹

$$W_x(v') \geq W_x(v'') - \sum_{x_n(d)=o} [v''_{n,o} - v'_{n,o}] = W(v', r).$$

Since we could construct v' for any $\varepsilon > 0$ (by Lemma A.4) and the welfare function $W_x(\cdot)$ is

²¹By construction, $v''_{n,o} \geq v'_{n,o}$ if $x_n(d) = o$ and $v''_{n,o} \leq v'_{n,o}$ if $x_n(d) \neq o$. We use the fact that the welfare is 1-Lipschitz in the first case and nondecreasing in the second case.

continuous, we have

$$W_x(v) \geq W(v, r);$$

since this is true for all $r \in R$, we have

$$W_x(v) \geq \max_{r \in R} \{W(v, r)\}. \quad (17)$$

Combining (15) and (17), we get

$$W_x(v) = \max_{r \in R} \{W(v, r)\},$$

so x is welfare-equivalent to a maximal-in-range algorithm with range R . \square

Proof of Theorem 2.6

Theorem 2.6. *For any weakly monotone algorithm x and any $\beta \in [0, 1]$, x is a β -approximation for investment if and only if x is a β -approximation for certain investment.*

The certain investment instances are a subset of the investment instances. Thus, if x is a β -approximation for investment, then x is a β -approximation for certain investment.

We now prove the other direction. Suppose we have some investment instance $(g, I, \nu_{-l}, \mathcal{A})$ and some algorithm x that is a β -approximation for certain investment. Going off the intuition described in the main text, we seek to construct state-dependent cost functions $\tilde{c} : S \rightarrow \mathbb{R}$ to make the *ex post* normalized utility from the investment constant (and equal to the *ex ante* normalized expected utility from that investment).

Now, the state-dependent cost function $\tilde{c} : S \rightarrow \mathbb{R}$ that gives constant utility in each state for investment (ν_l, c) is given by

$$\tilde{c}(s) \equiv u_l(\nu(s), \mathcal{A}(s), x) - \underbrace{\left(\left[\sum_{s'} g(s') u_l(\nu(s'), \mathcal{A}(s'), x) \right] - c \right)}_{\text{ex ante normalized utility from } (\nu_l, c)}.^{22}$$

Observe that $\sum_s g(s) \tilde{c}(s) = c$. By construction, the *ex post* normalized utility of choosing the investment with state-dependent-cost (ν_l, \tilde{c}) , i.e., $u_l(\nu(s), \mathcal{A}(s), x) - \tilde{c}(s)$, is in every state equal to the *ex ante* normalized utility of the original investment (ν_l, c) .

However, we need a status quo alternative so that the realization in each state is a valid investment instance, so we introduce a new investment option with 0 normalized utility and

²²Here we use the normalized utility notation introduced in the Preliminaries & Notation section at the start of the appendix.

0 cost. Formally, we define an investment option that yields a value in each state equal to the investor's threshold prices

$$\nu'_l(s) \equiv \tau_l(\nu_{-l}(s), \mathcal{A}(s), x).$$

Let I' consist of the investments modified to have state-dependent costs, as well as the status quo alternative, that is $I' \equiv \{(\nu_l, \tilde{c}) : ((\nu_l, c) \in I) \cup \{(\nu'_l, 0)\}$. For each state s , the tuple $(I'(s), \nu_{-l}(s), \mathcal{A}(s))$ is a certain investment instance. Since x is a β -approximation for certain investment, for any $(\nu_l, c) \in \text{BR}(x, g, I, \nu_{-l}, \mathcal{A})$ and any state s , we have

$$W_x(\nu(s), \mathcal{A}(s)) - \tilde{c}(s) \geq \beta \bar{W}^*(I'(s), \nu_{-l}(s), \mathcal{A}(s)).$$

Thus, for any $(\nu_l, c) \in \text{BR}(x, g, I, \nu_{-l}, \mathcal{A})$ such that $\bar{W}_x(I, \mathcal{A}) = [\sum_s g(s) W_x(\nu_l(s), \mathcal{A}(s))] - c$, we conclude that

$$\begin{aligned} \bar{W}_x(I, \nu_{-l}, \mathcal{A}) &= \left[\sum_s g(s) W_x(\nu(s), \mathcal{A}(s)) \right] - c \\ &= \sum_s g(s) [W_x(\nu(s), \mathcal{A}(s)) - \tilde{c}(s)] \\ &\geq \sum_s g(s) \beta \bar{W}^*(I'(s), \nu_{-l}(s), \mathcal{A}(s)) \\ &\geq \beta \bar{W}^*(I', \nu_{-l}, \mathcal{A}) \geq \beta \bar{W}^*(I, \nu_{-l}, \mathcal{A}), \end{aligned}$$

where the penultimate inequality holds because the expectation of the maximum is no less than the maximum of the expectation, and the final inequality holds because for all $(\nu_l, c) \in I$, we have $(\nu_l, \tilde{c}) \in I'$ with $\sum_s g(s) \tilde{c}(s) = c$.

Proof of Theorem 2.10

Theorem 2.10. *For any weakly monotone algorithm x and any $\beta \in [0, 1]$, if x has β -bounded confirming externalities, then x is a β -approximation for certain investment.*

We start with the following technical lemma.

Lemma A.5. *The optimal welfare function $W^*(v_n)$ is*

- *non-decreasing in v_n and*
- *1-Lipschitz in v_n*

in the sup norm.

Proof. This follows from Lemma A.3 because the optimal welfare function has zero externalities. \square

Next we state a technical condition that we will prove is equivalent to β -approximation for certain investment. The condition compares $W_x(v, A)$ minus one bidder's normalized utility to the optimal welfare when that bidder's value is set to the threshold vector.²³

Definition A.6. For some $\beta \in [0, 1]$, algorithm x is β -**pivotal** if for any allocation instance (v, A) and any bidder n , we have

$$W_x(v, A) - u_n(v, A, x) \geq \beta \underbrace{W^*(\tau_n(v_{-n}, A, x), v_{-n}, A)}_{\text{optimal welfare at } n\text{'s threshold vector}} .$$

Lemma A.7. For any weakly monotone algorithm x and any $\beta \in [0, 1]$, x is a β -approximation for certain investment if and only if x is β -pivotal.

Proof. Being a β -approximation for certain investment means the welfare for algorithm x must be lower bounded by β times the optimal welfare for any investment instance. However, many of these lower bounds turn out to be redundant, and our argument shows that the only relevant bounds are those in the β -pivotality condition.

Algorithm x is a β -approximation for certain investment if and only if for every set of investments I , any $(v_l, c) \in \text{BR}(x, I)$, and any $(v'_l, c') \in I$,

$$W_x(v_l) - c \geq \beta(W^*(v'_l) - c'). \quad (18)$$

If we restrict attention to investment sets of the form

$$I = \{(v_l, c), (v'_l, c'), (\tau_l, 0)\},$$

then we still obtain the bounds (18); however, we can rewrite the requirement that $(v_l, c) \in \text{BR}(x, I)$ as $u_l(v_l, c) \geq \max\{u_l(v'_l, c'), 0\}$. By formulating (18) as

$$W_x(v_l) - u_l(v_l) \geq \beta(W^*(v'_l) - u_l(v'_l)) - u_l(v_l, c) + \beta u_l(v'_l, c'),$$

we notice that the tightest bound occurs when $u_l(v_l, c) = u_l(v'_l, c') = 0$. Therefore algorithm x is a β -approximation for certain investment if and only if for any v_l , where

$$W_x(v_l) - u_l(v_l) \geq \max_{v'_l} \{\beta(W^*(v'_l) - u_l(v'_l))\}. \quad (19)$$

²³Here we use the normalized utility notation introduced in the Preliminaries & Notation section at the start of the appendix.

Since by Lemma A.5 we have

$$W^*(v'_l) - u_l(v'_l) = W^*(v'_l - u_l(v'_l)) \leq W^*(v'_l - (v'_l - \tau_l)) = W^*(\tau_l),$$

we know the maximum in (19) occurs at $v'_l = \tau_l$; plugging this in, (19) becomes

$$W_x(v_l) - u_l(v_l) \geq \beta W^*(\tau_l),$$

which is exactly the β -pivotality condition. Therefore algorithm x is a β -approximation for certain investment if and only if x is β -pivotal. \square

We are now ready to prove Theorem 2.10. Suppose that x has β -bounded confirming externalities. If the change from the threshold price τ_n to any value v_n is a confirming change, then x having β -bounded confirming changes would immediately imply that x is β -pivotal. Now, while it is not the case that the change from the threshold price τ_n to any value v_n is confirming, we show that it *is* true that for any value v_n , there exists a value v_n^ε arbitrarily close to τ such that the change from v_n^ε to v_n is a confirming change.

Formally, for any v_n and $\varepsilon \in (0, 1]$, we let

$$v_n^\varepsilon \equiv \varepsilon v_n + (1 - \varepsilon)\tau_n.$$

Now, because x is truthful, we have

$$\begin{aligned} x_n(v_n^\varepsilon) \in \operatorname{argmax}_o \{v_{n,o}^\varepsilon - \tau_{n,o}\} &= \operatorname{argmax}_o \{(\varepsilon v_{n,o} + (1 - \varepsilon)\tau_{n,o}) - \tau_{n,o}\} \\ &= \operatorname{argmax}_o \{\varepsilon[v_{n,o} - \tau_{n,o}]\} \\ &= \operatorname{argmax}_o \{v_{n,o} - \tau_{n,o}\}. \end{aligned}$$

It follows that for any outcome $o' \in O$,

$$v_{n,x_n(v_n^\varepsilon)} - \tau_{n,x_n(v_n^\varepsilon)} \geq v_{n,o'} - \tau_{n,o'}. \quad (20)$$

Thus, we see that the change from v_n^ε to v_n confirms $x_n(v_n^\varepsilon)$ because

$$\begin{aligned} v_{n,x_n(v_n^\varepsilon)} - v_{n,o'} &= [(v_{n,x_n(v_n^\varepsilon)} - \tau_{n,x_n(v_n^\varepsilon)}) - (v_{n,o'} - \tau_{n,o'})] + \tau_{n,x_n(v_n^\varepsilon)} - \tau_{n,o'} \\ &\geq \varepsilon[(v_{n,x_n(v_n^\varepsilon)} - \tau_{n,x_n(v_n^\varepsilon)}) - (v_{n,o'} - \tau_{n,o'})] + \tau_{n,x_n(v_n^\varepsilon)} - \tau_{n,o'} \\ &= v_{n,x_n(v_n^\varepsilon)}^\varepsilon - v_{n,o'}^\varepsilon, \end{aligned}$$

where the inequality follows from (20).

Since the change from v_n^ε to v_n confirms $x_n(v_n^\varepsilon)$ and x has β -bounded confirming externalities, we have

$$p_n(v_n) - p_n(v_n^\varepsilon) + \sum_{m \neq n} [v_m \cdot [x_m(v_n) - x_m(v_n^\varepsilon)]] = \mathcal{E}_x(v_n, v_n^\varepsilon) \geq \beta W^*(v_n^\varepsilon) - W_x(v_n^\varepsilon). \quad (21)$$

Using the definition of normalized utility u_n , the inequality (21) becomes

$$W_x(v_n) - W_x(v_n^\varepsilon) - u_n(v_n) + u_n(v_n^\varepsilon) \geq \beta W^*(v_n^\varepsilon) - W_x(v_n^\varepsilon). \quad (22)$$

Canceling $W_x(v_n^\varepsilon)$ from both sides of (22) and taking the limit as $\varepsilon \rightarrow 0$, we get

$$W_x(v_n) - u_n(v_n) \geq \beta W^*(\tau_n).$$

Thus, we see that x is β -pivotal, and by Lemma A.7, x is a β -approximation for certain investment.

Proof of Proposition 2.18

Proposition 2.18. *Suppose that $|O| = 2$, and let X be a collection of weakly monotone X CONE algorithms. If y is an algorithm that at each instance $(v, A) \in \Omega$ outputs a welfare-maximizing allocation from the collection $\{x(v, A)\}_{x \in X}$, then y is weakly monotone.*

Suppose that y and X satisfy the assumptions of Proposition 2.18. We want to prove that for any (v, A) and \tilde{v}_n ,

$$0 \leq [\tilde{v}_n - v_n] \cdot [y_n(\tilde{v}_n) - y_n(v_n)]. \quad (23)$$

If $y_n(\tilde{v}_n) = y_n(v_n)$ then (23) follows immediately. Suppose $y_n(\tilde{v}_n) \neq y_n(v_n)$. If the change from v_n to \tilde{v}_n confirms $y_n(\tilde{v}_n)$ then (23) follows immediately. Suppose it does not confirm $y_n(\tilde{v}_n)$. Then by $y_n(\tilde{v}_n) \neq y_n(v_n)$ and $|O| = 2$, the change from v_n to \tilde{v}_n confirms $y_n(v_n)$, and the change from \tilde{v}_n to v_n confirms $y_n(\tilde{v}_n)$.

Let us pick $x, \tilde{x} \in X$ such that $x(v_n) = y(v_n)$ and $\tilde{x}(\tilde{v}_n) = y(\tilde{v}_n)$. We have

$$\begin{aligned} v_n \cdot \tilde{x}_n(v_n) + \sum_{m \neq n} [v_m \cdot \tilde{x}_m(v_n)] &\leq v_n \cdot x_n(v_n) + \sum_{m \neq n} [v_m \cdot x_m(v_n)] \\ &\leq v_n \cdot x_n(\tilde{v}_n) + \sum_{m \neq n} [v_m \cdot x_m(\tilde{v}_n)], \end{aligned} \quad (24)$$

where the first inequality is by construction and the second inequality is by x XCONE and weakly monotone and Proposition 2.7. A symmetric argument yields

$$\begin{aligned} \tilde{v}_n \cdot x_n(\tilde{v}_n) + \sum_{m \neq n} [v_m \cdot x_m(\tilde{v}_n)] &\leq \tilde{v}_n \cdot \tilde{x}_n(\tilde{v}_n) + \sum_{m \neq n} [v_m \cdot \tilde{x}_m(\tilde{v}_n)] \\ &\leq \tilde{v}_n \cdot \tilde{x}_n(v_n) + \sum_{m \neq n} [v_m \cdot \tilde{x}_m(v_n)]. \end{aligned} \quad (25)$$

Adding inequalities (24) and (25) and canceling terms yields

$$0 \leq [\tilde{v}_n - v_n] \cdot [\tilde{x}_n(v_n) - x_n(\tilde{v}_n)]. \quad (26)$$

Since the change from v_n to \tilde{v}_n confirms $y_n(v_n) = x_n(v_n)$, and x_n is weakly monotone, we have

$$0 = [\tilde{v}_n - v_n] \cdot [x_n(\tilde{v}_n) - x_n(v_n)]. \quad (27)$$

Similarly, since the change from \tilde{v}_n to v_n confirms $y_n(\tilde{v}_n) = \tilde{x}_n(\tilde{v}_n)$, and \tilde{x}_n is weakly monotone, we have

$$0 = [\tilde{v}_n - v_n] \cdot [\tilde{x}_n(\tilde{v}_n) - \tilde{x}_n(v_n)]. \quad (28)$$

Adding (26), (27), and (28) yields

$$0 \leq [\tilde{v}_n - v_n] \cdot [\tilde{x}_n(\tilde{v}_n) - x_n(v_n)] = [\tilde{v}_n - v_n] \cdot [y_n(\tilde{v}_n) - y_n(v_n)],$$

as desired.

Proof of Proposition 3.8

Proposition 3.8. *For all $\delta > 0$, there exists $\varepsilon < \delta$ such that the BKV rule with parameter ε has an investment guarantee of 0.*

In the following family of examples, we consider $\varepsilon = \frac{1}{2^i}$ for $i \in \mathbb{N}$.

Let there be $|N| = 2^j$ bidders where the first bidder is the investor, the second bidder has a value of 2, and the remaining $|N| - 2$ bidders have a value of 1. The knapsack can fit either the first two bidders or everyone except for the second bidder.

Consider what happens when the investor has value 0. For the BKV rule, the allocation that maximizes the rounded values will have $\ell \leq i + j$. For $\ell \leq i + j$, the BKV rule will pack everyone except for the second bidder. If $\ell > i + j$, the last $|N| - 2$ bidders' values will be rounded to 0.

If the investor increases its value to 2^{i+j+2} , the allocation that maximizes the rounded

values will be for $\ell = i + j + 1$. If ℓ is smaller, the investor's value will be truncated and if ℓ is larger, the second bidder's value will be rounded down to 0. For $\ell = i + j + 1$, $\gamma^{\varepsilon, \ell} = \frac{1}{2}$ so all values will be rounded down to the nearest even integer. Since the last $|N| - 2$ bidders' values are rounded down to 0, the BKV rule will pack the first two bidders.

As shown above, the investor increasing its value from 0 to 2^{i+j+2} results in a confirming negative externality and therefore the BKV rule is not XCONE. If the set of investments is $\{(0, 0), (2^{i+j+2}, 2^{i+j+2})\}$, then the performance under investment is $\frac{2}{|N|-2}$ which goes to 0 as $|N|$ goes to ∞ .

Proof of Proposition 3.11

Proposition 3.11. *The allocation rule \check{x}^ε can be computed in $\text{poly}(\varepsilon^{-1}, |N|, \log(\max_n\{v_n\}))$ time.*

If $\frac{\varepsilon 2^\ell}{|N|} > \max_n\{v_n\}$, then every value rounds to 0, and by construction $x^*(v^{\varepsilon, \ell})$ packs no bidders and thus yields 0 welfare. Consequently, it suffices to compute $x^*(v^{\varepsilon, \ell})$ from $\ell = 0$ to $\ell = \lfloor \log(\varepsilon^{-1}|N|\max_n\{v_n\}) \rfloor + 1$ in order to find the best output from the collection $(\check{x}^{\varepsilon, \ell})_{\ell \in \mathbb{N}}$. Briest et al. (2005) proved that computing $x^*(v^{\varepsilon, \ell})$ in each step takes $\text{poly}(\varepsilon^{-1}, |N|, \log(\max_n\{v_n\}))$ time. Thus, we can compute \check{x}^ε in $\text{poly}(\varepsilon^{-1}, |N|, \log(\max_n\{v_n\}))$ time, which completes the proof of Proposition 3.11.

As an aside, we note that our proposal runs more slowly than the BKV FPTAS. The BKV FPTAS searches ℓ from $\lfloor \log(\max_n\{v_n\}) \rfloor - \lfloor \log((1-\varepsilon)^{-1}|N|) \rfloor - 1$ to $\lfloor \log(\max_n\{v_n\}) \rfloor$. Our FPTAS searches a larger range, from 0 to $\lfloor \log(\varepsilon^{-1}|N|\max_n\{v_n\}) \rfloor + 1$. This is because we are choosing the maximal allocation according to v instead of $v^{\varepsilon, \ell}$, and must search a larger range to find the relevant maximum.