Optimal Scaling Auctions: A Consumer Theory Deconstruction*

Allan Hernández-Chanto† — University of Queensland

Ernesto Rivera Mora‡ — University of Colorado, Boulder

This Version: October 14, 2025

First version: July 23, 2024

Abstract

We study scaling auctions in procurement settings with risk-averse sellers, showing that bidding behavior follows an auxiliary Hicksian demand in which scoring weights act as prices. As a result, changes in weights induce a pure substitution effect, providing a rationale for skewed bidding behavior observed in empirical studies. By using the Hicksian properties, we characterize the optimal scaling auction and show that it outperforms cash auctions if and only if marginal costs are sufficiently high. Additionally, we identify a hybrid auction format that outperforms both scaling and cash auctions.

JEL classification: C72, D44, D81, D82

Keywords: Utility equivalence; scoring auctions; scaling auctions; risk-averse sellers

*We are grateful for comments from Zachary Breig, Hector Chade, Yongmin Chen, Amanda Friedenberg, Marina Halac, Andreas Kleiner, Yao Luo, Alejandro Manelli, Claudio Mezzetti, Harry Pei, Daniel Rappoport, Larry Samuelson, and Philipp Strack, as well as audiences at Adolfo Ibañez University, Deakin University, Texas A&M University, The 2025 Australasian Economic Theory Workshop, The 2024 Texas Economic Theory Camp, The 2025 Stony Brook Economic Theory Conference, The 2025 Lisbon Meetings in Game Theory, The Paris Seminar on Game Theory, The Royal Melbourne Institute of Technology, The University of the Andes, The University of Auckland, The University of Colorado Boulder, The University of Graz, The University of New South Wales, The University of Otago, The University of Queensland, The University of Salamanca, The University of Texas at Dallas, The University of Vigo, and Yale University.

†Email: a.hernandezchanto@uq.edu.au

‡Email: ernesto.riveramora@colorado.edu

1 Introduction

Infrastructure investment is crucial for economies, as evinced by substantial public expenditure on projects such as the construction of bridges, highways, and public transit systems undertaken by local governments annually. Governments typically allocate projects to private construction firms through procurement auctions in which the winner is responsible for executing the project according to predetermined specifications and quality standards. These large projects often involve multiple tasks that are susceptible to unexpected changes during execution. For instance, soil conditions may differ from those in initial geological surveys, underground pipelines may require rerouting, or additional reinforcement may be needed for structural stability. These uncertainties create substantial cost volatility for both firms and governments, prompting public scrutiny over the optimal allocation of such financial risk.

The two predominant mechanisms used to allocate these projects are cash auctions (also known as fixed-price auctions) and scaling auctions (also known as unit-price auctions). In cash auctions, the government pays the winning firm a fixed monetary amount upfront, leaving the firm to bear all the risk entailed by cost variations. Conversely, scaling auctions—a form of scoring auction—allow firms to hedge against input-quantity risk. In this format, an engineer estimates the expected input quantities required for each task (for example, the number of barricades needed, or the liters of paint required for pavement markings). These estimated input quantities serve as the weights in the auction's scoring rule. Thus, the score associated with a given bidding vector is computed as the weighted average of the unit-price bids. The firm with the lowest score wins the project, but is paid their unit-price bids multiplied by the realized quantities, rather than the estimated quantities. The contingent nature of the payment makes scaling auctions particularly attractive, as it allows firms to hedge, transferring part of the risk to the government.

We investigate the optimal design of scaling auctions in procurement settings where a risk-neutral buyer (i.e., the government) faces a set of risk-averse sellers (i.e., the construction firms). We assume that sellers (i) have private heterogeneous fixed costs; (ii) have homogeneous marginal costs; and (iii) have a commonly known utility function that exhibits risk aversion but that does not exhibit wealth effects.[1] Our analysis yields four main results. First, in the methodological camp, we show that although sellers' equilibrium strategies vary with weights, their utility remains unchanged across all scaling auctions. This utility equivalence implies that any Pareto improvement from altering weights necessarily reduces the buyer's expected payment. Second, building on this result, we show that despite the vast range of weights the buyer could adopt, the common practice of setting weights equal to expected quantities minimizes the buyer's expected payment.[2] Third, we establish that optimal scaling auctions dominate cash auctions if and only if marginal costs are large relative to fixed costs. Finally, we show that a hybrid auction—in which sellers may request an upfront cash payment in addition to their per-unit bid—is strictly preferred by the buyer to both optimal pure scaling and cash auctions.

Our model incorporates heterogeneous, non-contractible fixed costs among sellers, a feature particularly relevant in the infrastructure sector. This sector brings together diverse sellers whose characteristics are shaped primarily by differences in fixed costs. In settings where scaling auctions apply, these costs vary substantially across sellers, reflecting factors such as capacity, organizational structure, and the opportunity costs of allocating resources to the project. For example, sellers with larger project backlogs often face higher expenses from maintaining a larger workforce or sustaining readiness across multiple projects. Likewise, sellers operating in multiple regions typically incur higher fixed costs due to resource duplication and differing regulatory requirements.

[1] We use the term *wealth effects* as used in the literature on risk aversion, in which the level of risk aversion is independent of wealth. We use *income effects* as employed in consumer theory, referring to changes in consumption choices from variations in income due to changes in prices.

[2] One might intuitively expect the optimal scoring rule should incorporate higher moments of the distribution: perhaps assigning greater weight to high-variance tasks could seemingly increase competition on low-variance tasks thereby reducing risk and expected payments. Our analysis shows this intuition is incorrect: using expected quantities is indeed optimal.

In addition, we consider sellers whose utility functions do not exhibit wealth effects. This assumption reflects the idea that, with perfect access to credit markets, firms can effectively manage liquidity constraints. Consequently, their degree of risk aversion remains unaffected by changes in wealth—leaving risk preferences as the primary driver of decision-making.

Finally, we assume that the buyer and sellers share common beliefs about project quantities. This assumption allows us to focus on how sellers allocate bids across tasks to hedge against the project's intrinsic risk. Importantly, unlike the empirical literature on scaling auctions, our model does not require normally distributed beliefs. While such an assumption can simplify analysis by enabling closed-form equilibrium solutions, our results hold for any distribution. Moreover, they remain valid in environments where sellers possess more information than the buyer—a natural setting, since sellers participate in these auctions repeatedly and become experts, while the government remains comparatively less informed.

Turning to the solution of the induced game, we follow Bolotnyy and Vasserman (2023) and Luo and Takahashi (2025) in deriving sellers' optimal bidding strategies as the solution of a two-stage maximization process: an initial stage in which they determine the optimal score to bid (i.e., "the outer loop") and a subsequent stage in which they optimize over the best way to achieve such a score given the weights of the scoring rule (i.e., "the inner loop"). Asker and Cantillon (2008) shows that the solution of this two-stage maximization process is equivalent to the solution of the original problem under fairly general conditions. Consequently, by using the indirect utility of each score, the seller's problem in the outer loop solves the trade-off between maximizing the score's indirect utility and maximizing the probability of winning.

We enhance this methodology by providing an equivalence between scaling auctions—in terms of strategies and induced payoffs—and a simplified game called the *auxiliary auction*. In the auxiliary auction, sellers compete by requesting "certainty equivalents" of the induced lottery over payments. We show that for any choice of weights in the scaling auction, the equilibrium in the associated auxiliary auction yields the same utility to each seller.[3] The buyer, on the other hand, prefers the vector of weights that delivers these equilibrium utility levels in the most efficient way, i.e., providing as much insurance as possible.

The sellers' utility equivalence between the scaling and auxiliary auctions allows us to apply a duality approach. First, we notice that the inner loop resembles the *utility maximization problem* in consumer theory. Thus, the optimal bid in this stage can be interpreted as the bidder's *Marshallian* demand in which the "vector of prices" corresponds to the vector of weights and the "income" corresponds to the intended score. Second, we show that the equilibrium bidding in the scaling auction can be obtained as a solution to a *score-minimization problem* constrained by the requirement to achieve a specific utility level in the auxiliary auction, similar to expenditure minimization in consumer theory. Hence, the equilibrium bid in the scaling auction can be interpreted as a *Hicksian demand*. Therefore, to determine the sellers' equilibrium bids we first translate the scaling auction into an auxiliary auction, calculate the interim expected utility for each seller's type in this simplified framework, and then translate the problem back into the scaling auction as a score minimization problem.

We use the duality approach to show that the optimal weights must be proportional to the expected quantities of tasks. This conclusion is grounded in three key observations present across all symmetric equilibria. First, sellers consistently obtain the same expected utility. Second, the buyer absorbs the risk from the sellers. Third, sellers determine their bids by minimizing their scores. As a result, when the weights are proportional to the expected quantities, the sellers' incentives align with those of the buyer, allowing the buyer to *delegate* the risk-minimization problem to the sellers. Our consumer-theory approach is not merely expository; it leverages duality in an innovative way to reduce complexity in scaling auctions—much as dual methods do

[3] Our online appendix fully characterizes necessary and sufficient conditions for this equivalence to hold in more general environments. A condition we call separability (which is satisfied here) serves as a sufficient condition.

in other complex environments, such as multi-unit object selling mechanisms.

Additionally, as a byproduct of our methodology, we obtain a version of the celebrated equivalence between Marshallian and Hicksian demands in the context of scaling auctions, which is novel to the literature. This allows us to formulate the *Law of Demand* for the vector of equilibrium bids. Specifically, if the weight of one task increases—holding all other weights constant—all the bids for this particular task decrease. In other words, the associated *income effect* is offset by the competitive nature of the auction, and thus, changes in equilibrium bids are fully attributed to a pure *substitution effect*.

Determining the optimal weights enables a comparison between the performance of cash auctions and optimal scaling auctions. To achieve this, we examine the *risk premium* faced by sellers in each auction format. In a cash auction, all risks are due to uncertainty over quantities and are independent of the sellers' fixed costs. Conversely, in a scaling auction, the uncertainty over quantities affects both the payment and the costs, which are reflected in the sellers' markups. Importantly, due to utility equivalence, these markups are determined by the auxiliary auction. As a result, the markups depend solely on the sellers' fixed costs and their informational rents in the induced game but not on their marginal costs. Identifying the source of risk in each format allows us to conclude that cash auctions minimize the buyer's expected payment when marginal costs are sufficiently low; otherwise, a scaling auction with optimal weights will result in a lower expected payment. This insight is consistent with current practices from departments of transportation that suggest project managers use cash auctions for unsophisticated projects and scaling auctions for complex large-scale projects.[4]

Building on the empirical results of Luo and Takahashi (2025), we also consider *hybrid auctions*, in which the winning seller's payment combines a lump-sum upfront

[4]For instance, the Florida Department of Transportation (FDoT) recommends using cash auctions for projects such as minor road widening, construction of sidewalks, and fencing; but recommends using scaling auctions for projects that involve subsoil earthwork, urban reconstruction, and for rehabilitation of movable bridges (see Luo and Takahashi, 2025).

payment with a contingent payment determined by the results of a scaling auction. When the lump-sum amount is exogenously set by the buyer, all of our results continue to hold, as this simply shifts the fixed costs by the lump-sum amount. However, when the lump-sum amount is endogenous, the scoring rule must account for both the bid requesting the lump-sum payment (i.e., the fixed-price bid) and the unit-price bids for each project task. We show that a hybrid auction with optimal weights allows sellers to fully hedge against the intrinsic risk of the project. This is because, in the associated portfolio problem, the fixed-price task serves as a risk-free asset, enabling sellers to cover fixed costs and informational rents without taking on additional risk. Therefore, under utility equivalence, an optimal hybrid auction generates a lower expected payment for the buyer compared to both cash and pure scaling auctions.

Finally, we discuss the effects of information acquisition. Specifically, while it has no impact on sellers, the buyer benefits both from allowing sellers to obtain more information and from acquiring information himself. The buyer can facilitate the former by encouraging site visits and achieve the latter through a two-stage process that elicits sellers' information. Allowing sellers to gather more information reduces auction risk, thereby benefiting the buyer.[5]

Related literature The theoretical literature has long examined the design of optimal scoring auctions, particularly in environments where sellers are risk-neutral and outcomes are deterministic after the auctioned project is allocated. In this line, Che (1993) and Awaya et al. (2022) determine the optimal scoring auction when sellers bid on both price and quality, with uni-dimensional types influencing their marginal costs. Asker and Cantillon (2008, 2010) expand this framework to environments where sellers' bids and types are arbitrarily multi-dimensional. They show that the optimal scoring auction can be derived by mapping each seller's multidimensional type into a

[5] Empirical evidence supports this notion. For example, Fok (2023) examines a procurement model using data from the Indiana Department of Transportation (INDoT). The counterfactual analysis suggests that preventing sellers from acquiring information could increase procurement costs by 1.62%, or about $37,000 per project. In INDoT's mechanism, sellers can anonymously ask questions about the cost structure, and responses are made publicly available to all sellers.

one-dimensional pseudo-type, allowing the use of standard techniques from Myerson (1981). More recently, Hanazono et al. (2024) studies price-per-quality-ratio scoring rules, a non-linear mechanism in which revenue equivalence breaks down due to the convexity of sellers' indirect utility in the score.

Drawing on these theoretical insights, there has been a growing empirical literature that applies and extends the principles of scoring auctions to various economic settings. Notable contributions include timber auctions (e.g. Athey and Levin, 2001), auctions for resolving failed banks (e.g., Allen, Clark, Hickman, and Richert, 2023), and auctions for environmental services (e.g., Aspelund and Russo, 2023).

In the realm of infrastructure procurement auctions, two recent studies stand out. Luo and Takahashi (2025) examines scaling and cash auctions used by the Florida Department of Transportation (FDoT), modeling risk-averse sellers with heterogeneous Gaussian beliefs, varying risk aversion, and endogenous entry. They show that scaling auctions reduce sellers' exposure to project risk, while cash auctions generate greater bid dispersion and discourage participation. Their model also allows for hybrid formats that combine lump-sum transfers with unit-price bids: sellers allocating more weight to non-lump-sum tasks bid more aggressively, reflecting the "insurance effect" of scaling auctions, though lump-sum items still mitigate risk exposure. Meanwhile, Bolotnyy and Vasserman (2023) studies data from the Massachusetts Department of Transportation (MDoT), where risk-averse sellers "skew" bids in anticipation of overruns and underruns. Scaling auctions again provide insurance, but sellers bear residual risk that commands a median premium of up to 14.5%. They further estimate that replacing scaling auctions with cash auctions would raise the cost of the median project by 42%.[6]

Our theoretical framework complements these recent empirical studies in two ways. First, it provides clear guidance on identifying the optimal auction format

[6]The authors also study an enhanced cash auction where sellers commit to fixed payments but renegotiate ex-post. This lowers the expected cost increase to 33.5%, yet still underperforms the baseline scaling auction.

based on project characteristics. Second, it offers a theoretical foundation for empirical regularities such as *bid skewing*, i.e., the tendency of sellers to submit low bids in some tasks (and high in others) to minimize their risk exposure and maximize profits.[7] For instance, Example 2 shows that bid skewing can emerge in optimal scaling auctions under risk-averse sellers and symmetric information between the buyer and sellers. Moreover, in this example, sellers optimize by "going short" (i.e., placing negative unit-price bids) in the task with the highest variance. This finding challenges conventional interpretations that view bid-skewing as evidence of either risk-neutrality or inaccurate scoring rules resulting from information asymmetry (Athey and Levin, 2001; Bajari, Houghton, and Tadelis, 2014; Bolotnyy and Vasserman, 2023).[8] This insight also has important implications for policy: bid skewing should be allowed, not prohibited. While some departments of transportation ban this practice, we find that bid skewing can actually help minimize sellers' risk exposure and consequently minimize the buyer's expected payment.

The risk-sharing feature embedded in scaling auctions naturally generalizes the classical results from the literature on optimal auctions under uncertainty (e.g., Riley and Samuelson, 1981; Maskin and Riley, 1984; Matthews, 1987). Within scaling auctions, sellers face not only the risk of losing the auctioned project but also uncertainty during the ex-post stage of the auction. This dual-layered uncertainty connects them to the literature on security-bid auctions, where the returns of the auctioned projects are stochastic, and financial instruments other than cash are used to secure payments (DeMarzo et al., 2005a). In this framework, Fioriti and Hernandez-Chanto (2022) shows that "steeper" security designs allow sellers to extract a greater surplus

[7] The literature typically attributes *"skewing behavior"* or *"unbalanced bids"* to asymmetric information between the buyer and the sellers. Under this interpretation, sellers increase unit prices for items they expect to exceed the buyer's quantity estimates while decreasing unit prices for items they expect to fall short of the buyer's estimates. Example 2 provides a novel rationale of this behavior.

[8] Athey and Levin (2001) notes that a risk-neutral seller's optimal strategy is to bid zero on overestimated items and bid all actual costs on underestimated ones. However, while zero unit price bids are observed, they are rare, which the authors attribute to risk aversion.

through the insurance they provide, further aligning this strand of research with the risk-sharing dynamics in scaling auctions.[9]

2 The environment

A group of $\mathcal{I} = \{1, 2, \ldots, I\}$ qualified sellers compete in a procurement auction for a contract to complete a project. The project encompasses a set of $\mathcal{T} = \{1, 2, \ldots, T\}$ tasks defined by the buyer at the outset of the bidding process, all of which must be completed for the successful execution of the project. Each task requires the use of different input quantities $\boldsymbol{q} = (q_t : t \in \mathcal{T})$ that are non-deterministic and over which sellers have a common prior belief $\mu \in \Delta(\mathbb{R}^T)$. The prior $\mu \in \Delta(\mathbb{R}^T)$ is *generic* in the sense that for each non-zero vector $\boldsymbol{v} \in \mathbb{R}^T$, $\mu\big(\{\boldsymbol{q} : \boldsymbol{q} \cdot \boldsymbol{v} = 0\}\big) < 1$.[10] Hence, no task can be deterministically described as a convex combination of the others. Moreover, the vector of realized quantities is independent of the outcome of the auction (e.g., the identity of the winning seller, the final payment, etc.).

Each seller i must incur the same marginal cost $c_{i,t} = c_t$ to procure the quantity q_t required for each task t. Additionally, prior to the project's execution, they must cover a fixed cost k_i. Sellers commonly know marginal costs, but the fixed cost k_i is seller i's private information and is drawn identically and independently from the commonly known absolutely continuous distribution F with density f. Thus, sellers' heterogeneity in efficiency is captured by differences in their fixed costs, which relate more to the organizational structure of the seller than to advantages in the inputs market. Therefore, in contrast to Bolotnyy and Vasserman (2023), we assume that firms face the same marginal costs but may differ in their fixed costs k_i.[11]

[9] Our Online Appendix, we show that both scaling auctions and security-bid auctions belong to a general class of scoring auctions under two layers of uncertainty, which can be analyzed within a unified theoretical framework of scoring auctions.

[10] This assumption is instrumental in obtaining strictly convex preferences over bids. The assumption is made without loss of generality: if μ is not generic—meaning one task is a linear combination of the others—redefine a new set of tasks $\mathcal{T}'$ such that $|\mathcal{T}'| < |\mathcal{T}|$ that has a non-generic prior belief.

[11] This cost structure extends the approach of Krasnokutskaya (2011), which is standard in the empirical literature for modeling unobserved heterogeneity, by specifying that the private cost component is bidder-specific rather than auction-specific.

Sellers are risk averse, with preferences represented by a strictly increasing and concave Bernoulli utility function $u : \mathbb{R} \to \mathbb{R}$ normalized so that $u(0) = 0$. We assume that the utility exhibits no wealth effects. This means that payoffs are evaluated independently of the seller's initial wealth.

The buyer uses either a scaling auction or a standard first-price auction to allocate the rights to complete the project among sellers.[12]

Discussion of the environment. Two key assumptions of our environment deserve more detailed discussion: null wealth effects and homogeneous marginal costs. We provide the rationale for each below.

First, the class of utility functions without wealth effects has been widely used in the finance literature for its analytical tractability. When combined with Gaussian beliefs, this class captures mean-variance preferences, a feature central to explaining phenomena such as the equity premium puzzle (Mehra and Prescott, 1985). Notably, this class is employed roughly five times more frequently than all other utility forms combined (Corner and Corner, 1995).[13]

Second, homogeneous marginal costs ensure that sellers hedge their variable costs consistently in response to weight changes. This assumption is reasonable in scaling auctions for public infrastructure (e.g., highways and bridges), in which production relies on standardized, capital-intensive inputs such as concrete and steel, as well as non-skilled labor, all of which are typically sourced from competitive markets. Consequently, marginal costs are largely driven by the competitive prices of these inputs and the wages of construction workers, while fixed costs vary substantially across sellers. Moreover, while there may be some specialized workers, such as designers or engineers, many of their contributions are one-time tasks whose costs can be interpreted

[12] Online Appendix B.5 provides an analysis of the second-price variants of these auctions.

[13] Our model retains the no-wealth-effects assumption familiar from these applications but dispenses with Gaussian beliefs, making it more general than mean-variance preferences and able to capture a broader class of environments.

as part of fixed costs.[14]

Under these assumptions, scaling auctions satisfy a condition we call *separability*, in which all types assign the same certainty equivalent value to each score. We show that separability implies utility equivalence among sellers.[15] As a result, these two conditions are sufficient for our methodology and ensure the model remains empirically relevant and analytically tractable. However, Section 7.3 shows that as the number of sellers approaches infinity, our results continue to hold even if these two assumptions are relaxed. In such environments, strong competition drives sellers' expected utility to zero, and the buyer's optimal weights are proportional to the expected quantities.

3 Scaling auctions

A *scaling auction* is characterized by a pair $\mathcal{A} = (\boldsymbol{w}, \mathcal{B})$ in which $\boldsymbol{w} = (w_t : t \in \mathcal{T}) \in \mathbb{R}^T_+$ is a non-zero vector of task-specific weights and $\mathcal{B} \subseteq \mathbb{R}^T$ is a convex and closed set of admissible bids. The scaling auction $\mathcal{A}$ is chosen at the beginning of the auction and is announced upfront by the buyer.[16] Each seller i must submit a per-unit bid for every task t featured in the project. These bids are submitted simultaneously and are kept confidential until the conclusion of the auction. For each seller i, the vector of bids $\boldsymbol{b}_i = (b_{i,t} : t \in \mathcal{T}) \in \mathcal{B}$ induces a score $s_i = \boldsymbol{b}_i \cdot \boldsymbol{w}$. The contract is

[14] The assumption of homogeneous marginal costs extends to various auction environments in which firms rely on standardized inputs. For example, in wholesale electricity markets, firms' marginal costs are largely determined by the competitive prices of inputs like natural gas and coal. As Hortaçsu and Puller (2008) discusses, standardized pricing mechanisms in commodity markets ensure that firms procure inputs at nearly identical prices, leading to convergence in their marginal production costs.

[15] Our online appendix studies a broader class of environments in which sellers' efficiency types influence multiple dimensions of heterogeneity, including marginal costs, fixed costs, and quantity distributions. There, we identify necessary and sufficient conditions for utility equivalence. A sufficient condition is separability: a scoring auction is separable if there exists a continuous, strictly increasing mapping $R : \mathbb{R} \to \mathbb{R}$ such that the expected utility of each type θ winning the project at score s is given by $u(R(s) - k(\theta))$, where $k(\theta)$ is the fixed cost for type θ.

[16] In many real auction settings $\mathcal{B} = \mathbb{R}^T_+$, meaning that only non-negative bidding vectors are accepted. We do not impose this restriction. In fact, Section 4.3 shows that for some beliefs μ, minimizing the expected payment may require negative bids for some tasks.

awarded to the seller with the *lowest score*, who is responsible for executing the entire project. Upon the project's completion, the winning seller is compensated based on her unit bid, $b_{i,t}$, multiplied by the actual quantity required by each task t, q_t. Thus, if seller i wins the auction, the total payment received from the buyer is $\boldsymbol{q} \cdot \boldsymbol{b}_i$ and her total ex-post monetary payoff is $\pi_i = \boldsymbol{q} \cdot (\boldsymbol{b}_i - \boldsymbol{c}) - k_i$. The latter accounts for the buyer's payment, the unit costs per task, and the fixed costs. Losing bidders do not pay or earn anything, regardless of their bid.

3.1 Bidding as a Marshallian demand

The realization of the actual quantities, $\boldsymbol{q} = (q_t : t \in \mathcal{T})$, used in each task of the project is unknown at the time of bidding. Hence, sellers need to compute the expected payoff induced by each potential bid. In doing so, they face two sources of uncertainty: the inherent uncertainty of winning, and the uncertainty associated with the stochastic nature of the actual quantities. Following Asker and Cantillon (2008), and more recently Bolotnyy and Vasserman (2023) and Luo and Takahashi (2025), we separate the seller's problem into two stages to account for both sources of uncertainty: (i) the *score bidding* problem and (ii) the *portfolio* problem. In the first stage, each seller i selects a score s_i based on her private fixed cost k_i, which determines their probability of winning the auction and imposes a constraint on their bidding decisions. Meanwhile, in the second stage, seller i chooses a vector of unit bids $\boldsymbol{b}_i$ that maximizes her interim expected utility, subject to the constraint $\boldsymbol{b}_i \cdot \boldsymbol{w} = s_i$. The restriction ensures that the chosen bidding vector induces the score selected in the first stage by the seller, given the vector of weights chosen by the buyer. As such, the second stage is equivalent to a utility-maximizing consumer's choice of a consumption bundle with a budget constraint entailed by prices $\boldsymbol{w}$ and an income s_i.

Interim and ex-ante expected utility Fix a type $k_i \in [\underline{k}, \overline{k}]$. The *interim expected utility* of a seller that wins the project with a bid $\boldsymbol{b}_i$ corresponds to:

$$U(\boldsymbol{b}_i \mid k_i) \equiv \mathbb{E}_{\mu}\Big[u(\boldsymbol{q} \cdot (\boldsymbol{b}_i - \boldsymbol{c}) - k_i)\Big]. \tag{1}$$

Hence, seller i's *ex-ante expected utility* from bidding $\boldsymbol{b}_i$ is given by

$$U(\boldsymbol{b}_i \mid k_i) \times \Pr\big[\boldsymbol{b}_i \cdot \boldsymbol{w} < s_j \text{ for all } j \neq i\big]. \tag{2}$$

The next result shows that $U(\cdot \mid k_i)$ satisfies various standard properties analogous to those in consumer theory.

Lemma 1. *For each fixed cost $k_i \in \mathbb{R}$, the interim expected utility $U(\cdot \mid k_i) : \mathbb{R}^T \to \mathbb{R}$ is continuous and strictly concave. Moreover, if the support of μ is bounded and contained in $\mathbb{R}^T_+$, then $U(\cdot \mid k_i)$ is continuously differentiable and strictly increasing in the component-wise linear order.*

To characterize sellers' equilibrium strategies in a scaling auction $\mathcal{A} = (\boldsymbol{w}, \mathcal{B})$, let $\mathcal{B}[s_i \mid \boldsymbol{w}] \equiv \{\boldsymbol{b}_i \in \mathcal{B} : \boldsymbol{b}_i \cdot \boldsymbol{w} \leqslant s_i\}$ denote the lower contour set of the score s_i. Hence, a score $s_i \in \mathbb{R}$ is *achievable* for $\mathcal{A}$ if $\mathcal{B}[s_i \mid \boldsymbol{w}]$ is non-empty. Fix a feasible score s_i. Notice that, given the seller's strictly convex preferences, the portfolio problem to achieve the score s_i has a unique maximizer, denoted as

$$\boldsymbol{b}_i(s_i \mid k_i, \mathcal{A}) = \arg \max_{\boldsymbol{b}_i \in \mathcal{B}[s_i \mid \boldsymbol{w}]} \big\{U(\boldsymbol{b}_i \mid k_i)\big\}. \tag{3}$$

The solution to the maximization problem, $\boldsymbol{b}_i(s_i \mid k_i, \mathcal{A})$, is called seller i's *Marshallian bidding demand.* The following result shows that the Marshallian bidding demand is invariant with respect to the fixed cost k_i.

Lemma 2. *Fix a scaling auction $\mathcal{A}$. For each feasible score s_i and each pair $k_i, k_i' \in [\underline{k}, \overline{k}]$, $\boldsymbol{b}_i(s_i \mid k_i, \mathcal{A}) = \boldsymbol{b}_i(s_i \mid k_i', \mathcal{A})$.*

Lemma 2 follows directly from the absence of wealth effects in the seller's utility function. This feature ensures that all seller types have identical indifference curves over bids, which decouples the bidding problem from the sellers' private information. Consequently, the Marshallian bidding demand does not depend on fixed costs. Given this independence, we simplify our notation by omitting k_i and write the seller's optimal bid as $\boldsymbol{b}_i(s_i \mid \mathcal{A})$. The following example illustrates the bidding demand in a setting with two tasks.

Example 1. *Assume that there are two tasks. The quantities* (q_1, q_2) *are drawn from a normal distribution with means* $(3, 2)$*, unitary variance, and zero correlation.*[17] *There are two competing sellers, each with fixed costs independently drawn from a uniform distribution in* $[0, 1]$*. The marginal cost of each task is* $\frac{1}{2}$*. Sellers have a utility function* $u(\pi) = 1 - \exp(-3\,\pi)$*. Figure 1 illustrates the seller's indifference curves and the Marshallian demand for each score and weight* $\boldsymbol{w} = (3, 2)$*. The green dashed line describes the bids that induce the same score. (By Lemma 1 and Lemma 2, the utility curves and the demand are independent of the realized fixed cost* k_i*.)*

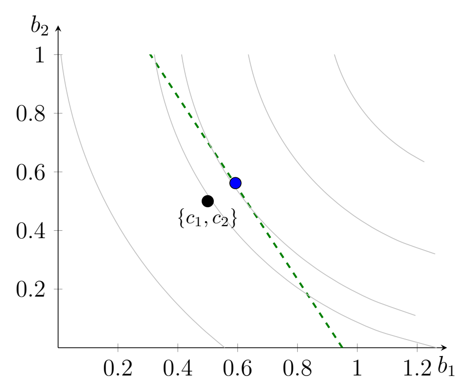

Figure 1: Seller's Marshallian demand for score s_i in auction $\mathcal{A}$.

[17] While inadequate in some circumstances, Gaussian beliefs provide a simple closed-form solution for the portfolio problem. (See Bolotnyy and Vasserman (2023).)

Indirect interim utility and certainty equivalence Fix $\mathcal{A} = (\boldsymbol{w}, \mathcal{B})$ and write

$$V(s_i \mid k_i, \mathcal{A}) \equiv U\big(\boldsymbol{b}_i(s_i \mid \mathcal{A}) \mid k_i, \mathcal{A}\big),$$

for the indirect utility associated with the maximization problem.

Fix a feasible score s_i of $\mathcal{A}$ and write $\mathrm{CE}(s_i \mid \mathcal{A}) \in \mathbb{R}$ for the value that solves

$$u\big(\mathrm{CE}(s_i \mid \mathcal{A})\big) = \mathbb{E}_\mu\big[u\big(\boldsymbol{q} \cdot (\boldsymbol{b}_i(s_i \mid \mathcal{A}) - \boldsymbol{c})\big)\big]. \tag{4}$$

That is, $\mathrm{CE}(s_i \mid \mathcal{A})$ is the unique monetary payment that makes the seller i indifferent between receiving a (deterministic) payment of $\mathrm{CE}(s_i \mid \mathcal{A})$ and receiving the stochastic payment of the lottery $\boldsymbol{q} \cdot (\boldsymbol{b}_i(s_i \mid \mathcal{A}) - \boldsymbol{c})$.[18] Note, since the utility function has no wealth effects, it follows that for each $k_i \in [\underline{k}, \overline{k}]$,

$$u\big(\mathrm{CE}(s_i \mid \mathcal{A}) - k_i\big) = \mathbb{E}_\mu\big[u\big(\boldsymbol{q} \cdot (\boldsymbol{b}_i(s_i \mid \mathcal{A}) - \boldsymbol{c}) - k_i\big)\big] = V\big(s_i \mid k_i, \mathcal{A}\big).$$

That is, for each seller i with fixed cost k_i, the utility from the certainty equivalent, after subtracting the fixed cost, equals the indirect utility obtained in the scaling auction. (This observation is central to the definition of an auxiliary auction, as discussed in Section 4.)

Now, as it is customary in the design of any allocation mechanism, the buyer must provide incentives to sellers to participate in the procurement auction. The following definition warrants that all sellers have incentives to participate.

Definition 1. *Call an auction $\mathcal{A}$ rich if for each fixed cost $k_i \in [\underline{k}, \overline{k}]$, there is some feasible score s_i for $\mathcal{A}$ such that $\mathrm{CE}(s_i \mid \mathcal{A}) = k_i$.*

In a rich auction $\mathcal{A}$, the bidding set $\mathcal{B}$ is sufficiently broad to ensure two key properties. First, participation in the auction is individually rational for all types of firms: for every type k_i, there exists a sufficiently high score s_i such that $\mathrm{CE}(s_i \mid$

[18] Since $u(\cdot)$ is strictly increasing, there is a unique real value that satisfies Equation (4).

$\mathcal{A}) \geqslant k_i$. Second, with enough competition, it is possible to extract all surplus from sellers: for each type k_i, there exists a sufficiently low score s'_i such that the certainty equivalent $\mathrm{CE}(s'_i \mid \mathcal{A}) \leqslant k_i$.[19] The buyer's objective is to identify a rich scaling auction $\mathcal{A}$ that minimizes the expected cost of procuring the completion of the project.[20]

3.2 Sellers' equilibrium scores

We now analyze the equilibrium bidding behavior in the scaling auction by considering the solution to the Marshallian bidding problem described above. Notice that seller i's ex-ante expected utility can be separated into two distinct components: the probability of winning and the interim expected utility conditional on winning. Given the distribution of scores of other participants, the probability of seller i winning the auction is determined solely by her score s_i. Consequently, all unit bids that adds up to the same score will result in an identical probability of winning. Meanwhile, the interim utility of seller i, conditional on winning, depends exclusively on the specific unit bids sellers submit, disregarding the bids of any other participants.

Following the tradition of the literature, we focus on the analysis of symmetric equilibria, where all participants adopt identical strategies based on their private information and the structure of the auction. Denoting $\varphi_{\mathcal{A}} : [\underline{k}, \bar{k}] \to \mathbb{R}$ as the strictly monotone equilibrium bidding function of a scaling auction $\mathcal{A}$, we can write the expected utility of bidding a score s_i as

$$\mathrm{EU}(s_i \mid k_i, \mathcal{A}) := V\big(s_i \mid k_i, \mathcal{A}\big) \cdot \left[1 - F\left(\varphi_{\mathcal{A}}^{-1}(s_i)\right)\right]^{I-1}.$$

[19] If the support of μ contains only vectors with non-negative quantities, then any auction with $\mathcal{B} = \mathbb{R}_+^T$ is rich. If beliefs are Gaussian, the mapping $\mathrm{CE}(\cdot \mid \mathcal{A})$ is bounded above by some value $\overline{\mathrm{CE}} \in \mathbb{R}$, and a rich auction does not exist if $\overline{\mathrm{CE}} < \bar{k}$. If $\overline{\mathrm{CE}} \geqslant \bar{k}$, then each scaling auction with $\mathcal{B} = \mathbb{R}_+^T$ is rich.

[20] If the auction does not provide incentives for all types to participate, the project will not be executed with probability one. Likewise, if the auction cannot extract all surplus from sellers under high competition, the buyer leaves money on the table and fails to minimize expenditure. Section 7.2 analyzes optimal non-rich auctions and shows that they can reduce expected payments, but only at the cost of delivering the project with probability strictly less than one.

Because fixed costs are independently and identically drawn from an absolutely continuous and atomless distribution F, the probability of having two different sellers with the same equilibrium score is zero.[21]

Definition 2. *Fix a scaling auction $\mathcal{A} = (\boldsymbol{w}, \mathcal{B})$. The score bidding function $\varphi_{\mathcal{A}}$ is a symmetric Bayesian equilibrium of $\mathcal{A}$ if for each $i \in \mathcal{I}$, each $k_i \in [\underline{k}, \overline{k}]$, and each achievable score $s_i \in \mathbb{R}$,*

$$\mathrm{EU}(\varphi_{\mathcal{A}}(k_i) \mid k_i, \mathcal{A}) \geqslant \mathrm{EU}(s_i \mid k_i, \mathcal{A}).$$

Characterizing the equilibrium often requires solving a differential equation that incorporates the integral defining $V(s_i \mid k_i, \mathcal{A})$. The conventional method typically involves placing restrictions on μ to derive a closed-form solution for $V(s_i \mid k_i, \mathcal{A})$. For example, Bolotnyy and Vasserman (2023) assume that μ follows a normal distribution to analyze the equilibrium strategy. In contrast, our approach does not rely on distributional restrictions to obtain a simple closed-form solution. Instead, we reformulate the strategic interaction as a task-free equivalent game that enables us to characterize the equilibrium for a broad class of beliefs.

4 Hicksian decomposition of scaling auctions

In this section, we show that the strategic behavior in the scoring auction can be equivalently described through a simple game of incomplete information, which we call *the auxiliary auction*. Importantly, the auxiliary auction omits the provision of quantities $\boldsymbol{q}$, yet preserves the strategic incentives inherent in the original scaling auction. This equivalence makes it possible to use the celebrated duality result in consumer theory to obtain the solution to the "portfolio problem" in the scaling auction through a score-minimization problem. The importance of this result is that

[21] The paper refers only to payoffs on the equilibrium path. Hence, we avoid writing the dependence of $\mathrm{EU}(s_i \mid k_i, \mathcal{A})$ on the equilibrium strategy $\varphi_{\mathcal{A}}$ for simplicity.

we can link the buyer's quest to minimize the expected payment to the seller's score minimization problem.

The auxiliary auction In the auxiliary auction, each seller has a fixed cost k_i, which is private information. Each seller i bids for a payment $r_i \in \mathbb{R}$ that they want to receive with certainty. The seller who bids the lowest r_i wins the auction, pays a fixed cost k_i, and receives r_i as a payment from the buyer. All sellers that lose the auction receive a zero monetary value. Thus, in its design, the auxiliary auction resembles a standard first-price auction over a project with non-stochastic value.

The utility of a winning seller with cost k_i which bids r_i is $u(r_i - k_i)$. Thus, if the sellers symmetrically bid according to a strictly increasing mapping $\psi : [\underline{k}, \overline{k}] \to \mathbb{R}$, then the ex-ante expected utility of such a seller bidding r_i, is

$$\mathrm{EU}_{\mathrm{aux}}(r_i \mid k_i) := u(r_i - k_i) \cdot [1 - F(\psi^{-1}(r_i))]^{I-1}.$$

Definition 3. *The mapping $\psi : [\underline{k}, \overline{k}] \to \mathbb{R}$ is a symmetric Bayesian equilibrium of the auxiliary auction if for each $i \in \mathcal{I}$, $\mathrm{EU}_{\mathrm{aux}}(\psi(k_i) \mid k_i) \geqslant \mathrm{EU}_{\mathrm{aux}}(r_i \mid k_i)$, for all $k_i \in [\underline{k}, \overline{k}]$, and all $r_i \in \mathbb{R}$.*

Lemma 3. *[(Krishna, 2009)] There exists a unique symmetric equilibrium $\psi : [\underline{k}, \overline{k}] \to \mathbb{R}$ of the auxiliary auction, which is characterized by the following differential equation:*

$$\psi'(k_i) = \frac{(I-1)\, f(k_i)\, u(\psi(k_i) - k_i)}{\big(1 - F(k_i)\big)\, u'(\psi(k_i) - k_i)} \tag{5}$$

and the boundary condition $\psi(\bar{k}) = \bar{k}$.[22]

[22]Comparative statics are easy to analyze within the solution to this differential equation. For instance, increasing competition (increasing I) decreases the bidding strategy of all seller types. Moreover, if u is a CARA utility function with risk aversion parameter γ, the bidding strategy decreases in γ. These comparative statics extend to equilibrium behavior of scaling auctions.

4.1 Sellers' utility equivalence

We show that any equilibrium of the original scaling auction can be written in terms of the unique equilibrium of the auxiliary auction. This is important since the latter does not depend on the particular rules of $\mathcal{A} = (\boldsymbol{w}, \mathcal{B})$, and thus, provides a way to circumvent the fact that $\boldsymbol{w}$ and $\mathcal{B}$ affect sellers' equilibrium behavior in a complex way. In fact, although no close solution can be written in the original scaling auction, the analysis of the auxiliary auction allows for comparative statics across different weighting vectors.

The key to the utility-equivalence result is that, regardless of the sellers' fixed costs, all sellers' types agree about the certainty equivalent value $\mathrm{CE}(s_i \mid \mathcal{A})$ of each bidding score s_i. This property, which we call *separability*, enables our consumer theory approach to characterize sellers' bidding behavior across different scaling auctions.[23]

Proposition 1. *Write $\psi : [\underline{k}, \bar{k}] \to \mathbb{R}$ for the unique symmetric equilibrium of the auxiliary auction. If the scaling auction $\mathcal{A} = (\boldsymbol{w}, \mathcal{B})$ is rich, then there exists a unique symmetric equilibrium $\varphi_{\mathcal{A}} : [\underline{k}, \bar{k}] \to \mathbb{R}$ which satisfies $\mathrm{CE}(\varphi_{\mathcal{A}}(k_i) \mid \mathcal{A}) = \psi(k_i)$. Moreover,*

$$U\big(\varphi_{\mathcal{A}}(k_i) \mid k_i, \mathcal{A}\big) = u\big(\psi(k_i) - k_i\big) \quad \text{and} \quad \mathrm{EU}\big(\varphi_{\mathcal{A}}(k_i) \mid k_i, \mathcal{A}\big) = \mathrm{EU}_{\mathrm{aux}}\big(\psi(k_i) \mid k_i\big).$$

Proposition 1 provides a complete characterization of the bidding behavior for each rich scaling auction. Notably, the two equalities above imply that the informational rents of the scaling auction are equivalent to the informational rents of the auxiliary auction. Therefore, the sellers' expected utilities are contingent upon the distribution F of fixed costs but remain unaffected by the prior beliefs μ, the marginal costs $\boldsymbol{c}$, the vector of weights $\boldsymbol{w}$, and the bidding set $\mathcal{B}$.

[23] Our Online Appendix analyzes separability within more general environments.

Corollary 1 (Utility equivalence). *Fix two rich scaling auctions* $\mathcal{A} = (\boldsymbol{w}, \mathcal{B})$ *and* $\mathcal{A}' = (\boldsymbol{w}', \mathcal{B}')$. *If* $\varphi_{\mathcal{A}} : [\underline{k}, \bar{k}] \to \mathbb{R}$ *is a bidding equilibrium for* $\mathcal{A}$ *and* $\varphi_{\mathcal{A}'} : [\underline{k}, \bar{k}] \to \mathbb{R}$ *is a bidding equilibrium for* $\mathcal{A}'$, *then* $\mathrm{EU}(\varphi_{\mathcal{A}}(k_i) \mid k_i, \mathcal{A}) = \mathrm{EU}(\varphi_{\mathcal{A}'}(k_i) \mid k_i, \mathcal{A}')$ *for all* $k_i \in [\underline{k}, \overline{k}]$.

The utility-equivalence result implies that all sellers are indifferent across rich scaling auctions, but the buyer is not. Different auctions involve varying levels of risk, leading firms to hedge with different expected payments. It is therefore in the buyer's interest to choose the auction that offers the greatest insurance.

A central novelty of our framework is that it allows for Pareto improvements in auction design. Because sellers' expected utilities are identical across all rich scaling auctions, the buyer can reduce his expected payment by selecting the scoring rule that provides the greatest insurance, without lowering sellers' welfare. This stands in sharp contrast to standard auction environments, where lower buyer payments often come at the expense of sellers' informational rents.

4.2 Equilibrium bidding as a Hicksian demand

This section uses utility equivalence to show that the sellers' optimal bids can be obtained as the solution of a dual score minimization problem subject to the constraint that a type k_i seller achieves a fixed utility level $\bar{u}_i$. This problem is analogous to the well-known expenditure minimization problem in consumer theory. Here, the score function is the dot product of a basket of "goods" $\boldsymbol{b} = (b_t : t \in \mathcal{T})$ and a vector of "prices" $\boldsymbol{w} = (w_t : t \in \mathcal{T})$.

Fix a scaling auction $\mathcal{A} = (\boldsymbol{w}, \mathcal{B})$, a fixed cost $k_i \in [\underline{k}, \overline{k}]$, a utility level $\bar{u}_i \in \mathbb{R}$, and write $B(\bar{u}_i \mid k_i, \mathcal{A}) := \left\{\boldsymbol{b}_i \in \mathcal{B} : U(\boldsymbol{b}_i \mid k_i) = \bar{u}_i\right\}$ for the bidding vectors in $\mathcal{B}$ that achieve the utility level $\bar{u}_i$. Call $\bar{u}_i$ *feasible* for $\mathcal{A}$ and k_i if $B(\bar{u}_i \mid k_i, \mathcal{A})$ is non-empty. Since $U(\cdot \mid k_i)$ is strictly convex, for each feasible $\bar{u}_i$ there exist a unique minimizer

$$\boldsymbol{h}_i(\bar{u}_i \mid k_i, \mathcal{A}) \equiv \arg \min_{\boldsymbol{b} \in B(\bar{u}_i \mid k_i, \mathcal{A})} \boldsymbol{b} \cdot \boldsymbol{w}.$$

We call $\boldsymbol{h}_i(\bar{u}_i \mid k_i, \mathcal{A})$, the *Hicksian bidding demand.*[24]

Proposition 2. *(Duality) Fix a scaling auction $\mathcal{A} = (\boldsymbol{w}, \mathcal{B})$ and a fixed cost $k_i \in [\underline{k}, \bar{k}]$. If $\bar{u}_i = u\big(\psi(k_i) - k_i\big)$, then $\boldsymbol{b}_i\big(\varphi_{\mathcal{A}}(k_i) \mid \mathcal{A}\big) = \boldsymbol{h}_i(\bar{u}_i \mid k_i, \mathcal{A})$. Moreover, if $\hat{\boldsymbol{b}}_i \in B(\bar{u}_i \mid k_i, \mathcal{A})$, then*

$$\boldsymbol{b}_i\big(\varphi_{\mathcal{A}}(k_i) \mid \mathcal{A}\big) \cdot \boldsymbol{w} \leqslant \hat{\boldsymbol{b}}_i \cdot \boldsymbol{w}.$$

Proof. For a formal proof see Proposition 3.E.1 Mas-Colell et al. (1995). □

Our Hicksian-style approach can be interpreted as an application of duality in auctions, resembling the techniques used in multi-object mechanism design Manelli and Vincent (2006). While multi-unit mechanism design considers independent valuations over multiple objects, our scaling auction involves a single large project partitioned into interdependent tasks. In both cases, the methodological principle is the same: by transforming the original problem into a dual formulation—here via the auxiliary auction, which incorporates project uncertainty into a certainty-equivalent utility—we decouple the two layers of risk, leaving only the risk of losing the auction. Sellers then minimize their score conditional on this target utility, yielding a tractable and exact characterization of equilibrium bidding strategies. Framing the approach through consumer theory and Hicksian demand is therefore not merely a formal convenience; it encapsulates a duality-based solution that handles the complex stochastic allocation environment of scaling auctions.

Substitution and income effects Notice that a change in the vector $\boldsymbol{w}$ has two effects in the equilibrium bid $\boldsymbol{b}_i(\varphi_{\mathcal{A}}(k_i) \mid \mathcal{A})$: it changes the associated relative cost of the bids and it changes the equilibrium score $\varphi_{\mathcal{A}}(k_i)$.[25] Proposition 2 shows that the sum of the two effects leads the sellers to face a pure *substitution effect* and a null *income effect.* That is, by the utility-equivalence result, the income effect associated with the changes of $\boldsymbol{w}$ is offset by the associated change in the equilibrium score

[24]Note, in contrast to the Marshallian demand, the Hicksian demand depends on k_i.

[25]Classic results such as Roy's identity and Shephard's lemma, extend to this auction setting.

$\varphi_{\mathcal{A}}(k_i)$. Intuitively, the competitive nature of the auction environment erases the income effect because there is only one project and the most efficient seller always obtains the project. Hence, the equilibrium score strategy $\varphi_{\mathcal{A}}(k_i)$ is shifted in a way that the sellers solely face a substitution effect. Interestingly, only the buyer reaps the benefits (or incurs the costs) of the substitution effect, as sellers obtain the same utility across all rich scaling auctions. Therefore, we obtain a version of the classical Law of Demand in consumer theory to capture the effects of changing the weight vector $\boldsymbol{w}$ in terms of the seller's *skewed bidding behavior*. Specifically, a bid $b_{t,i}$ weakly decreases if—ceteris paribus—the weight w_t increases. That is, increasing the weight w_t makes bidding on task t unappealing for the sake of winning the auction. Thus, all sellers decrease their bids for task t and increase their bids on other tasks.

4.3 Optimal scaling auctions

For any fixed rich scaling auction $\mathcal{A}$, we call $\mathcal{A}$ *optimal* if it minimizes the buyer's expected cost across all rich scaling auctions. This section employs the utility-equivalence result to characterize the class of optimal auctions.

Proposition 2 implies that the seller's equilibrium bidding vector always minimizes the scoring rule among all vectors of bids that yield the seller the same utility. Hence, by the utility equivalence, if weights are proportional to $\mathbb{E}_\mu[\boldsymbol{q}]$, then the sellers choose a bidding vector that minimizes the associated expected payment. The following Theorem illustrates this point.

Theorem 1. *Fix a rich scaling auction $\mathcal{A} = (\boldsymbol{w}, \mathcal{B})$. If $\mathcal{B} = \mathbb{R}^T$ and $\boldsymbol{w} = \lambda \cdot \mathbb{E}_\mu[\boldsymbol{q}]$ for some $\lambda > 0$, then $\mathcal{A}$ is optimal.*

Theorem 1 shows that optimal weights should be proportional to expected quantities. As the left panel of Figure 2 illustrates, any deviation from these weights distorts how sellers substitute bids across tasks, leading to higher expected payments from the buyer. When weights are proportional to expected quantities, the sellers' score-minimization problem aligns perfectly with the buyer's cost-minimization objective.

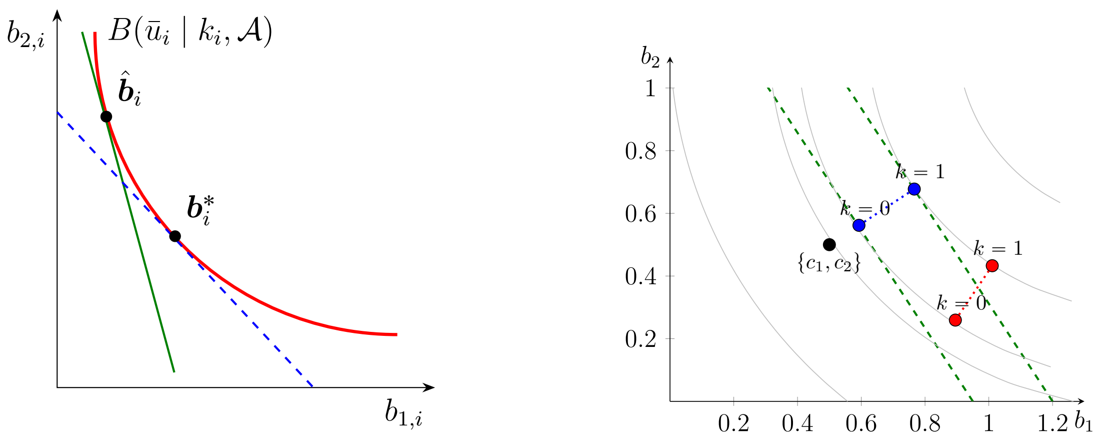

Figure 2: Left panel: The bid $\boldsymbol{b}^*$ represents the equilibrium bid under the optimal weights $\boldsymbol{w} = \lambda \mathbb{E}_\mu[\boldsymbol{q}]$. In contrast, the bid $\hat{\boldsymbol{b}}$ corresponds to the equilibrium bids under the suboptimal weights $\hat{\boldsymbol{w}}$. The solid green line denotes the iso-score curve for the suboptimal weights, while the dashed blue line represents the iso-score (and, incidentally, the iso-cost) curve under the optimal weights. **Right panel:** Bidding behavior and expected costs for an optimal and a suboptimal vector for Example 1.

This alignment effectively delegates the buyer's optimization problem to the sellers, as their individual incentives lead to choices that minimize the buyer's expected costs.

Consider the environment described in Example 1 and the (rich) scaling auctions $\mathcal{A} = (\boldsymbol{w}, \mathcal{B})$ and $\mathcal{A}' = (\boldsymbol{w}', \mathcal{B}')$, where $\mathcal{B} = \mathcal{B}' = \mathbb{R}^T$, $\boldsymbol{w} = (3,2) = \mathbb{E}_\mu[\boldsymbol{q}]$, and $\boldsymbol{w}' = (2,3) \neq \mathbb{E}_\mu[\boldsymbol{q}]$. Theorem 1 implies that auction $\mathcal{A}$ is optimal. The right panel of Figure 2 illustrates how auction $\mathcal{A}'$ induces a higher expected payment from the buyer. The gray lines represent the sellers' indifference curves in the bid space (note that these curves are independent of fixed costs). The black circle represents the firms' marginal costs. The blue dots depict, respectively, the bidding behavior of the least-efficient (i.e., $k = 1$) and most-efficient (i.e., $k = 0$) sellers under auction $\mathcal{A}$, while the blue dotted line shows the bids of intermediate sellers. Similarly, the red dots and red dotted line represent bidding behavior under auction $\mathcal{A}'$. The green dashed lines are the buyer's iso-cost curves.

Notice that the suboptimal weight vector $(2,3)$ distorts bidding incentives by making task 1 relatively more attractive. Following the "Law of Demand," sellers respond by shifting their bids toward task 1, leading them to bid below marginal cost for task 2. This bidding distortion results in auction $\mathcal{A}'$ generating a higher expected payment compared to the optimal auction $\mathcal{A}$.

Skewing behavior The empirical literature has documented that in multiple scenarios where scaling auctions are conducted, sellers have a tendency to skew bids of some tasks towards zero. The following example provides a rationale for this skewing bids even in the case of *risk aversion*, *symmetric information*, and *optimal weights*. Moreover, this skewing behavior is so extreme that the optimal scaling auction induces agents to bid negative bids in one of the tasks.

Example 2. *In this example, marginal costs are zero (the results also hold for small positive marginal costs). Two competing sellers have fixed costs uniformly distributed in* $[0, 2]$ *and utility function* $u(\pi) = 1 - \exp(-\pi)$. *There are only two tasks and the belief* μ *assigns a uniform probability of* $\frac{1}{2}$ *to* $(q_1, q_2) = (4, 0)$ *and* $(q_1, q_2) = (6, 10)$. *Note that expected quantities are* $\mathbb{E}_\mu[\boldsymbol{q}] = (5, 5)$.

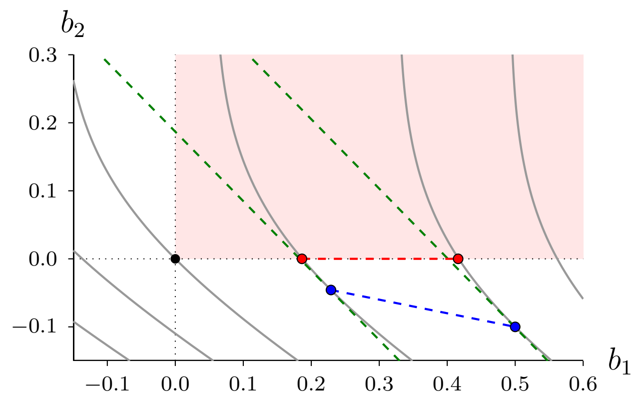

Figure 3: Bidding behavior under auctions $\mathcal{A}$ (respectively $\mathcal{A}'$) in blue (respectively in red). The optimal behavior under non-negative bids requires sellers to place a zero bid in the second task—i.e., induces a corner solution.

Fix $\boldsymbol{w} = (5, 5)$ *and consider two scaling auctions:* $\mathcal{A} = (\boldsymbol{w}, \mathcal{B})$ *and* $\mathcal{A}' = (\boldsymbol{w}, \mathcal{B}')$, *where* $\mathcal{B} = \mathbb{R}^T$ *and* $\mathcal{B}' = \mathbb{R}^T_+$. *So, both auctions use the optimal weights, but the second restricts bids to be non-negative. Theorem 1 implies that* $\mathcal{A}$ *is optimal. We show that*

$\mathcal{A}'$ is not optimal. Notice, absent any bidding restrictions, the demand associated with the weight vector $\boldsymbol{w}$ satisfies: $b_{i,1}(s_i \mid \mathcal{A}) = -5\, b_{i,2}(s_i \mid \mathcal{A}) \geqslant 0$ for each score $s_i > 0$. So, each seller i goes "short" against the risky task $t = 2$, fully hedging against risk; i.e., the payment $\boldsymbol{b}_i(s_i \mid \mathcal{A}) \cdot \boldsymbol{q}$ is constant.[26] *Observe that the restriction $\mathcal{B}' = \mathbb{R}^T_+$ in $\mathcal{A}'$ precludes these bids, forcing sellers into corner solutions that increase both risk exposure and the buyer's expected payment. Figure 3 shows the equilibrium bidding of $\mathcal{A}$ (in blue) and $\mathcal{A}'$ (in red).*

5 Cash auctions

In this section, we compare the expected payments and sellers' utilities of the scaling auction with those of a cash auction. In a cash auction, each seller i submits the total payment $r_i \in \mathbb{R}_+$ she would like to receive from the buyer upfront for the entire project. The seller offering the lowest bid wins the auction, receives the specified amount, and executes the project, incurring her own marginal and fixed costs.

5.1 Sellers' utility equivalence

To characterize the strategic behavior of this auction, it is convenient to introduce the certainty equivalent associated to the lottery of the tasks' costs. Fix the vector of marginal costs $\boldsymbol{c} = (c_t : t \in \mathcal{T})$ and denote $z \in \mathbb{R}$ as the unique constant such that $u(-z) = \mathbb{E}_\mu\big[u\,(-\boldsymbol{q} \cdot \boldsymbol{c})\,\big]$. That is, z represents the *certain equivalent cost* that each seller is willing to pay to avoid the uncertain burden of the tasks' costs.

Because there are no wealth effects, the expected utility for seller i upon winning the auction—after receiving the fixed payment r_i and paying the fixed cost k_i—is given by $u(r_i - k_i - z)$. Therefore, the utility of winning is equivalent to the utility in a variant of the auxiliary auction in which each firm faces a total cost of $k_i + z$ instead of only k_i. This observation leads to the lemma below.

[26] Under the consumer-theory interpretation, the bid associated with task 2 is an inferior good; i.e., the demand for $b_{i,2}$ decreases as the "income" s_i increases.

Lemma 4. *If $\zeta : \mathbb{R} \to \mathbb{R}$ is a symmetric equilibrium of the cash auction and ψ is a symmetric equilibrium of the auxiliary auction, then $\zeta(k_i) = \psi(k_i) + z$.*

Proposition 3 (Utility Equivalence)**.** *In each rich scaling auction $\mathcal{A}$, as well as in the auxiliary and the cash auction, a seller with fixed cost k_i expects to attain the same utility.*

5.2 Buyer's cost comparison

We now compare the buyer's costs incurred in both the cash auction and the scaling auction. Since sellers receive the same expected utility in both formats, the buyer absorbs the risk faced by sellers. Thus, it is in the buyer's best interest to design an auction that provides insurance to sellers in the most efficient way.

For each $\boldsymbol{\lambda} = (\lambda_t : t \in \mathcal{T}) \in \mathbb{R}^T$, write $\mathrm{RP}(\boldsymbol{\lambda}) \in \mathbb{R}$ for the unique constant such that $u\big(\mathbb{E}_\mu [\boldsymbol{\lambda} \cdot \boldsymbol{q}] - \mathrm{RP}(\boldsymbol{\lambda})\big) = \mathbb{E}_\mu\big[u(\boldsymbol{\lambda} \cdot \boldsymbol{q})\big]$. Call $\mathrm{RP}(\boldsymbol{\lambda}) \in \mathbb{R}$ the **risk premium** of lottery $\boldsymbol{\lambda} \cdot \boldsymbol{q}$ given a belief μ.[27] Thus, the risk premium is the certain payment that the sellers are willing to pay to avoid the uncertainty induced by the lottery.

Because utility equivalence holds, the auction that minimizes the buyer's expected payment is precisely the one that minimizes the risk premium borne by the sellers. Hence, the key is to compare the risk premium that each auction induces on the winner of the project. In a cash auction, the winning seller receives a fixed cash revenue and faces only uncertainty about the tasks' costs. Hence, the risk premium induced by the cash auction is $\mathrm{RP}(-\boldsymbol{c})$. By contrast, in a scaling auction $\mathcal{A}$, the winner faces uncertainty of both the task's cost and payment provided by the buyer.

For any scaling auction $\mathcal{A}$, the seller i's equilibrium markup for task t under $\mathcal{A}$ is defined as $m_{i,t}(k_i \mid \mathcal{A}) \equiv \boldsymbol{b}_{i,t}(\varphi_{\mathcal{A}}(k_i) \mid \mathcal{A}) - c_t$. Observe that, given the vector of markups $\boldsymbol{m}_i(k_i \mid \mathcal{A}) := (m_{i,t}(k_i \mid \mathcal{A}) : t \in \mathcal{T})$, the risk premium induced by the optimal-scaling auction is $\mathrm{RP}(\boldsymbol{m}_i(k_i \mid \mathcal{A}))$.

[27] Notice, for each vector $\boldsymbol{\lambda}$, the left-hand side is strictly decreasing in $\mathrm{RP}(\boldsymbol{\lambda})$. Hence, there exists a unique constant that satisfies this equation.

Lemma 5. *Fix $k_i \in [\underline{k}, \overline{k}]$ and write $\overline{u}_i = u\big(\psi(k_i) - k_i\big)$. If $\mathcal{A}$ is optimal, then the vector of equilibrium markups $\boldsymbol{m}_i(k_i \mid \mathcal{A})$ is the unique vector $\hat{\boldsymbol{m}} \in \mathbb{R}^T$ that minimizes $\mathbb{E}_\mu[\boldsymbol{q}] \cdot \hat{\boldsymbol{m}}$ subject to the constraint $\mathbb{E}_\mu\left[u\left(\boldsymbol{q} \cdot \hat{\boldsymbol{m}} - k_i\right)\right] = \overline{u}_i$.*

This result shows that the markups of the optimal scaling auction are determined solely by the informational rents of the auxiliary auction. Moreover, since the auxiliary auction solely depends on the distribution of sellers' fixed costs, the markups of any optimal scaling auction do not depend on the vector of marginal costs. In other words, if marginal costs increase from $\boldsymbol{c}$ to $\boldsymbol{c} + \boldsymbol{\epsilon}$, each seller i will respond by increasing their bid from $\boldsymbol{b}_i$ to $\boldsymbol{b}_i + \boldsymbol{\epsilon}$, leaving markups constant. Building on this result, Theorem 2 establishes a simple condition for comparing the expected payment between cash and optimal scaling auctions.

Theorem 2. *Let $\mathcal{A}$ be an optimal scaling auction. For each distribution of fixed costs F, there exists some constant $C_F > 0$ such that the expected payment of the cash auction exceeds the expected payment of $\mathcal{A}$ if and only if $\mathrm{RP}\left(-\boldsymbol{c}\right) \geqslant C_F$.*

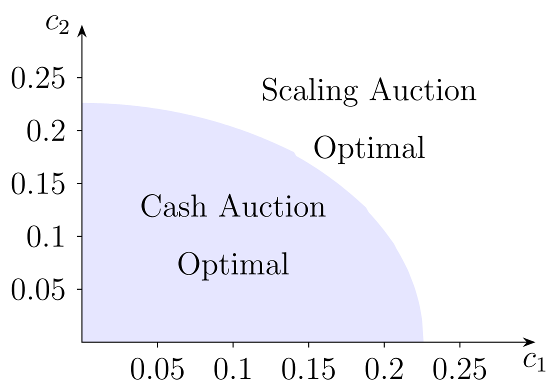

Figure 4: Optimal format in terms of marginal costs (Example 1).

Theorem 2 shows that the cash auction dominates scaling auctions if and only if the marginal costs $\boldsymbol{c}$ are sufficiently low relative to fixed costs. Figure 4 illustrates Theorem 2 in the context of Example 1, considering various levels of marginal costs while keeping all other parameters constant.

The proof of Theorem 2 consists of two main components. The first part states that the auction format minimizing the expected payment also minimizes the risk premium faced by the winning seller. Due to the equivalence in seller utilities, the risk-neutral buyer absorbs the entirety of sellers' risk. Consequently, the optimal auction provides greater insurance to the sellers.

The second part characterizes the kind of projects in which the scaling auction is optimal. It indicates that the scaling auction yields the best results in environments where marginal costs are relatively high. This conclusion stems from two key observations: Firstly, the risk premium of a cash auction depends solely on the costs of the tasks and remains independent of the sellers' fixed costs. Secondly, the risk premium of a scaling auction, C_F, depends solely on the distribution of fixed costs and is independent of the costs of the tasks. (See Lemma 5.) As a result, the key argument lies in comparing marginal costs with fixed costs.

As publicized by several departments of transportation, some projects are capital-intensive and have high marginal costs, while others are labor-intensive and have low marginal costs. Examples of the former are light bridges, high-speed railways, and other projects that require expensive and specialized materials (such as high-strength steel, carbon fiber-reinforced polymers, titanium, pre-stressed concrete, etc). Theorem 2 states that for such projects the scaling auction (under optimal weights) yields lower expected payments. Conversely, examples of low-marginal cost projects are development projects, rural infrastructure, and general maintenance (projects using gravel, plain concrete, clay bricks, maintenance goods, etc). In such projects, cash auctions induce a lower expected payment. This distinction is consistent with the empirical findings of Luo and Takahashi (2025).

6 Hybrid auctions

We will now analyze hybrid auctions. Hybrid auctions are formats in which, in addition to the payment $\boldsymbol{q} \cdot \boldsymbol{b}_i$, sellers receive a lump sum payment that can be either fixed or endogenous to the auction.

6.1 Hybrid auctions with fixed lump sum

First we consider the variant with a fixed lump-sum payment. In this case, the buyer selects the weight vector $\boldsymbol{w}$, the bidding set $\mathcal{B}$, and an exogenous lump sum payment $\ell \in \mathbb{R}$. Then, each seller i selects their bidding vector $\boldsymbol{b}_i \in \mathcal{B}$. The seller with the lowest score wins and receives the contingent payment $\boldsymbol{q} \cdot \boldsymbol{b}$, along with the lump sum payment.

The auction with a fixed lump sum transfer can be linked to a standard scaling auction by considering a shift in the distribution of fixed costs. In this alternate environment, each seller faces a fixed cost of $k_i - \ell$. Thus, the distribution of nominal fixed costs F and the lump sum ℓ induce a shift in the distribution $\widetilde{F}$ of real fixed costs over the set $[\underline{k} - \ell, \overline{k} - \ell]$. Therefore, because sellers' utility equivalence holds for each $\ell \in \mathbb{R}$, all the results of Section 4 continue to hold.

6.2 Hybrid auctions with endogenous lump sum

Building on the empirical results in Luo and Takahashi (2025), we consider an auction format in which the lump sum received by the winning seller is endogenously determined by sellers.[28] In this setting, a hybrid scaling auction is defined as a pair $\hat{\mathcal{A}} = (\hat{\boldsymbol{w}}, \hat{\mathcal{B}})$ with a vector of weights $\hat{\boldsymbol{w}} = (w_t \mid t \in \{0\} \cup \mathcal{T}) \in \mathbb{R}_+^{T+1}$. Here, $w_0 \in \mathbb{R}_+$ represents the weight of the lump sum, and w_t corresponds to the weight of task t. The bidding set $\hat{\mathcal{B}} \subseteq \mathbb{R}^{T+1}$ consists of the bidding vectors available to the sellers. Each seller i selects a monetary payment $\hat{b}_{0,i} \in \mathbb{R}$ that they wish to receive with certainty. The seller's bidding vector is denoted as $\hat{\boldsymbol{b}} \in \hat{\mathcal{B}}$, and the score rule is given by $s_i = \hat{\boldsymbol{w}} \cdot \hat{\boldsymbol{b}}_i$.

We define an extended set of tasks $\hat{\mathcal{T}} = \{0, \ldots, T\}$ and a random vector $\hat{\boldsymbol{q}} = (\hat{q}_t : t \in \hat{\mathcal{T}})$, where $\hat{q}_0 = 1$ and $\hat{q}_t = q_t$ for each $t > 0$. Furthermore, we define $\hat{\boldsymbol{c}} \equiv (0, c_1, \ldots, c_T)$, where the dummy task $t = 0$ has no marginal costs. Additionally,

[28] The amount requested by sellers can be either lower or higher than their fixed costs.

let $\hat{\mu} \in \Delta(\mathbb{R}^{T+1})$ represent the "extension" of the prior belief μ on $\mathbb{R}^{T+1}$, describing the probability distribution of $\hat{\boldsymbol{q}}$.[29] In this modified environment, the results from Sections 4 and 5 continue to hold.

Definition 4. *Fix a hybrid scaling auction* $\hat{\mathcal{A}} = (\hat{\boldsymbol{w}}, \hat{\mathcal{B}})$ *and let* $\hat{\boldsymbol{m}}(\cdot \mid \hat{\mathcal{A}}) : [\underline{k}, \overline{k}] \to \mathbb{R}^{T+1}$ *be its equilibrium markups. We say that* $\hat{\mathcal{A}}$ *provides full insurance if* $\mathrm{RP}\big(\hat{\boldsymbol{m}}(k_i \mid \hat{\mathcal{A}})\big) = 0$ *for each type* k_i*.*

Theorem 3. *Fix a hybrid scaling auction* $\hat{\mathcal{A}} = (\hat{\boldsymbol{w}}, \hat{\mathcal{B}})$*. If* $\hat{\boldsymbol{w}} = \big(1, \mathbb{E}[q_1], \ldots, \mathbb{E}_\mu[q_T]\big)$ *and* $\mathbb{R}^T_+ \subseteq \hat{\mathcal{B}}$*, then* $(\hat{\boldsymbol{w}}, \hat{\mathcal{B}})$ *is rich, optimal, and provides full insurance. Moreover,* $\hat{\mathcal{A}}$ *induces a lower payment than a cash auction and any rich non-hybrid scaling auction.*

Theorem 3 shows that when hybrid auctions use expected quantities as weights, they fully eliminate sellers' exposure to quantity risk.[30] By utility equivalence, hybrid scaling auctions then dominate both cash and non-hybrid scaling formats, lowering expected payments relative to either. This result complements Luo and Takahashi (2025) by providing a theoretical foundation: interpreting the lump sum as a "risk-free asset" allows sellers to fully hedge quantity risk and, through utility equivalence, strictly reduce expected payments, thereby rationalizing their empirical findings.

7 Extensions

This section explores several extensions of our theoretical framework. We examine how the analysis changes when agents acquire information; when strong bidding restrictions are imposed; when sellers exhibit heterogeneous characteristics; and when a second-price auction format is implemented.

[29] We assume that the belief $\hat{\mu}$ is generic. That is, no task in the extended environment is a linear combination of the other tasks.

[30] This implies that in hybrid auctions with endogenous lump-sum payments and optimal weights (proportional to expected quantities), all equilibrium bidding prices are non-negative. Hence, it is sufficient to restrict the bidding space to $\mathbb{R}^T_+$. Since negative prices may appear *non-credible* or *economically implausible*, this feature provides an additional advantage of hybrid auctions.

7.1 The value of information

This model assumes that both the sellers and the buyer share a common prior belief μ about the vector of tasks $\boldsymbol{q}$. Nevertheless, the empirical literature has noted that sellers are typically better informed than the buyer. Our model can be readily extended to a framework in which all sellers observe a public signal that the buyer does not observe.

Notice, if all sellers observe the same public signal, the same argument applies and the sellers' utility is still determined by the equilibrium of the auxiliary auction. Thus, utility equivalence holds and (public) information about tasks has no impact on the sellers' welfare. Intuitively, the public signal has two opposing effects that cancel each other out. The first is a standard *informativeness effect* as described by Blackwell (1951, 1953). The public signal enables sellers to reduce quantity uncertainty. Consequently, holding the behavior of other sellers fixed, winning the auction becomes more attractive. The second effect is an indirect *competitiveness effect*. As the reward for winning increases, sellers respond by bidding more aggressively, lowering their bids and thereby reducing their own welfare. These two effects exactly offset each other, as required by utility equivalence.

Although information is neutral for the sellers, the buyer can achieve better outcomes by either (i) allowing sellers to obtain more information or (ii) acquiring more information himself. So, overall, the buyer benefits from acquiring and revealing as much information as possible. To see the effect of (i), observe that, by allowing sellers to obtain more information, the risk premium associated with the auction diminishes. Hence, the benefit is transferred from the sellers to the buyer due to utility equivalence.[31] As for the effect of (ii), acquiring more information about tasks allows the buyer to make more informed decisions regarding the vector of weights $\boldsymbol{w}$. However, the benefits of information acquisition are confined to specific dimensions of information. Specifically, the magnitude of $\boldsymbol{q}$ becomes irrelevant since setting a vector

[31] Observe that this effect also holds for the cash auction.

of weights $\lambda \boldsymbol{w}$ is equivalent to setting $\boldsymbol{w}$. What becomes crucial for the buyer is to obtain information about the *direction* of $\boldsymbol{q}$. So, the gain comes from discerning the types of tasks that are more likely to be required, rather than estimating the project's overall size. For instance, in the construction of a bridge, determining the ratio of close substitutes (e.g., "water-proof" concrete vs. "low-shrinkage" concrete) is more important than estimating the bridge's total size. In the extreme case in which all project inputs are perfect complements, allowing the buyer to acquire additional information does not reduce his costs.

There are several strategies the buyer can employ to obtain these informational benefits. For instance, to achieve (i), the buyer may allow (and encourage) sellers to visit the prospect site of the project as a way to gather crucial information—e.g., soil and ground conditions, topography, existing infrastructure, and site accessibility. Moreover, the buyer can achieve (ii) by using a two-stage process to elicit the sellers' information. In the first stage, the buyer sets a vector of arbitrary weights and elicits sellers' homogeneous beliefs. In the second stage, the buyer recalibrates the weights using the elicited beliefs.

7.2 Exclusion restrictions

This section explores non-rich auctions, i.e., auctions in which the bidding space is small, precluding some seller types from participating.[32] Excluding certain seller types has three notable consequences. First, it directly breaks utility equivalence. Second, by reducing competition, it alters the bidding behavior of remaining sellers. Third, it allows the possibility of not completing the project, an outcome the buyer seeks to avoid.

Let $D > 0$ denote the exogenous social benefit of the project, which the buyer forfeits if no seller completes it. To solve for the optimal non-rich scaling auction, we

[32] For instance, the MDoT retains the authority to reject heavily skewed bids Bolotnyy and Vasserman (2023). Bajari, Houghton, and Tadelis (2014) document that the California Department of Transportation often rejects bids that are highly skewed.

rely on the tools developed in Section 4. Consider a scaling auction $\mathcal{A}$, and define $R_{\mathcal{A}} \equiv \{\mathrm{CE}(s_i \mid \mathcal{A}) : s_i \text{ is feasible under } \mathcal{A}\}$ as the set of constant equivalent payments feasible under $\mathcal{A}$. Let $\kappa_{\mathcal{A}} \equiv \max\{[\underline{k}, \overline{k}] \cap R_{\mathcal{A}}\}$ be the cutoff type for $\mathcal{A}$, i.e., the highest fixed cost at which a seller can participate. Accordingly, only seller types in $[\underline{k}, \kappa_{\mathcal{A}}]$ participate in $\mathcal{A}$.

Each auction can be associated with a variant of the auxiliary auction. Specifically, fix a convex set $R \subset \mathbb{R}$. An auxiliary auction with a restricted bidding set R modifies the standard auxiliary auction by restricting sellers to select bids from R.

Lemma 6. *Fix a scaling auction $\mathcal{A} = (\boldsymbol{w}, \mathcal{B})$. For each symmetric equilibrium $\varphi_{\mathcal{A}} : [\underline{k}, \kappa_{\mathcal{A}}] \to \mathcal{B}$ of $\mathcal{A}$, there exists a symmetric equilibrium $\psi : [\underline{k}, \kappa_{\mathcal{A}}] \to R_{\mathcal{A}}$ of the auxiliary auction with restricted bidding set $R_{\mathcal{A}}$ that satisfies $\mathrm{CE}(\varphi_{\mathcal{A}}(k_i) \mid \mathcal{A}) = \psi(k_i)$. Moreover, for each $k_i \in [\underline{k}, \kappa_{\mathcal{A}}]$,*

$$U\big(\varphi_{\mathcal{A}}(k_i) \mid k_i, \mathcal{A}\big) = u\big(\psi(k_i) - k_i\big) \quad \textit{and} \quad \mathrm{EU}\big(\varphi_{\mathcal{A}}(k_i) \mid k_i, \mathcal{A}\big) = \mathrm{EU}_{\mathrm{aux}}\big(\psi(k_i) \mid k_i\big).$$

Lemma 6 relates the equilibria between $\mathcal{A} = (\boldsymbol{w}, \mathcal{B})$ and the auxiliary auction with bidding set $R_{\mathcal{A}}$. Observe, restricting the set of bids to $\mathcal{B}$ restricts the set $R_{\mathcal{A}}$, thereby decreasing the set of sellers willing to participate. This induces the standard trade-off in procurement: restricting the set of sellers lowers the bids of the participating sellers but decreases the probability of finding a seller who completes the project.

As the social benefit D increases, the optimal design should accommodate a wider range of seller types to minimize the probability of project failure. Consequently, the optimal bidding set $\mathcal{B}$ is contingent on the magnitude of the social benefit D, which is generally larger for higher values of D. Notably, it remains optimal to employ a weight vector $\boldsymbol{w} = \mathbb{E}_{\mu}[\boldsymbol{q}]$, which incentivizes bidding in a manner that minimizes the buyer's expected payment, regardless of the set of firms that participate.

Remark 1. *Reducing the sellers' bids due to bidding restrictions requires that the buyer has a strong level of commitment power. The buyer would need to convince* *sellers that the project would be permanently canceled in case no seller participates.*

7.3 Heterogeneous characteristics and wealth effects

A natural extension of our framework considers sellers who have wealth effects and are heterogeneous not only in their fixed costs but also in their marginal costs and levels of risk aversion. We capture this rich heterogeneity by equipping each seller with a private type $\theta \in \Theta$. Each type θ maps to a fixed cost $k(\theta)$, a vector of marginal costs $\boldsymbol{c}(\theta) = (c_t(\theta) : t \in \mathcal{T})$. Each type θ has a utility function u_θ (wealth effects are allowed. The type of each seller is independently and identically distributed from a cumulative distribution function G in the support Θ.

Allowing a rich heterogeneous environment has significant implications. First, the choice of vector weights $\boldsymbol{w}$ becomes crucial in determining sellers' competitiveness, leading different types to win under different weight vectors. Thus, the utility equivalence result and the Hicksian properties of sellers' equilibrium behavior no longer hold. Nevertheless, as the number of sellers approaches infinity, utility equivalence is restored. In this limit, competition drives the expected utility of all sellers to zero.[33]

In this limiting case the choice of the vector of weights remains critical. Let $U(\boldsymbol{b} \mid \theta)$ denote the expected utility of a seller with type θ upon winning the auction with bid $\boldsymbol{b}$, and define $B_0(\theta) := \{\boldsymbol{b} \in \mathbb{R}^T : U(\boldsymbol{b} \mid \theta) = 0\}$ as the set of bids yielding zero expected utility for type θ. As we show below, this implies that the winning seller effectively minimizes the score subject to $\boldsymbol{b} \in B_0(\theta)$.

Proposition 4. *Assume there is a continuum of sellers and that the support of μ is bounded and contained in $\mathbb{R}^T_+$. If $\boldsymbol{w} = \lambda \mathbb{E}_\mu[\boldsymbol{q}]$ for some $\lambda > 0$ and $\mathcal{B} = \mathbb{R}^T$, then $\mathcal{A} = (\boldsymbol{w}, \mathcal{B})$ is an optimal scaling auction.*

[33] This outcome resembles the equilibrium of a Bertrand competition with firms facing multidimensional marginal costs. Notably, the result does not require imposing a lattice structure on the action space or supermodularity conditions, as is typically assumed in models such as Vives (1999).

8 Conclusion

Departments of transportation worldwide commonly use scaling auctions to procure public infrastructure projects. While these auctions share similar designs across governments, several key features remain underexplored in the academic literature. In particular, practices such as weighting bids based on government engineer estimates, prohibiting small (i.e., penny) bids, and permitting combinations of cash and unit-price bids lack thorough theoretical examination. Considering the substantial financial stakes involved, even minor modifications to these design elements could significantly reduce the expected costs of such projects.

We develop a model that offers formal answers to these questions in a general environment. Our framework allows for non-Gaussian seller beliefs and incorporates heterogeneous fixed costs into sellers' cost structures. We show that the prevailing practice of weighting bids proportionally to estimated quantities achieves optimality within the class of scaling auctions. However, we find that common restrictions such as prohibiting small "penny" bids (or negative bids) can constrain sellers' ability to hedge against quantity uncertainty, potentially increasing the buyer's expected costs. Moreover, hybrid auctions that combine cash and unit-price bids allow sellers to better diversify risk, thereby achieving the lowest expected payment for the buyer.

To characterize behavior in scaling auctions we introduce a novel methodology that uses a consumer theory approach to determine how sellers behave in equilibrium. Specifically, we map the scaling auction into the auxiliary auction in which sellers ask for sure payoffs instead of taking part in the auction's lottery. We then find each seller type's indirect utility in the auxiliary auction, map this value back to the scaling auction, and solve for the equilibrium bid as a score minimization problem. In this way, the Hicksian properties of bidding behavior allow us to disentangle how the environment's different features affect the auctions' efficiency and cost.

A key step in our methodology is establishing the sellers' utility equivalence between the scaling auction and the auxiliary auction. Our approach extends to a broader class of scoring auctions, either in (1) environments with many sellers or (2) environments without wealth effects, and with homogeneous marginal costs and input distributions. Our Online Appendix 8 provides a complete characterization of the necessary and sufficient conditions for applying this approach to this general class of scoring auctions that includes other important formats such as security bid auctions.

Importantly, our framework does not explicitly model sellers' moral hazard or the possibility of default. The latter aspect makes small (penny) bids unappealing, as inefficient types would have an incentive to win the auction with artificially small bids, collect partial payments, and then default on the project's delivery.

Our framework provides a foundation for analyzing these issues in richer environments. Laboratory experiments complement this work by enabling controlled study of counterfactual scenarios in complex auction settings. Since observational data are scarce, experiments offer unique insight into behavioral and informational mechanisms that remain hidden in field data.

References

Allen, J., R. Clark, B. Hickman, and E. Richert (2023). Resolving failed banks: Uncertainty, multiple bidding and auction design. *Review of Economic Studies*, rdad062.

Asker, J. and E. Cantillon (2008). Properties of scoring auctions. *The RAND Journal of Economics 39*(1), 69–85.

Asker, J. and E. Cantillon (2010). Procurement when price and quality matter. *The Rand Journal of Economics 41*(1), 1–34.

Aspelund, K. M. and A. Russo (2023). Additionality and Asymmetric Information in Environmental Markets: Evidence from Conservation.

Athey, S. and J. Levin (2001). Information and competition in us forest service timber auctions. *Journal of Political economy 109*(2), 375–417.

Awaya, Y., N. Fujiwara, and M. Szabo (2022, May). Quality and price in scoring auctions. Technical report. Technical Report.

Bajari, P., S. Houghton, and S. Tadelis (2014). Bidding for incomplete contracts: An empirical analysis of adaptation costs. *American economic review 104*(4), 1288–1319.

Blackwell, D. (1951). Comparison of experiments. In *Proceedings of the second Berkeley symposium on mathematical statistics and probability*, Volume 2, pp. 93–103. University of California Press.

Blackwell, D. (1953). Equivalent comparisons of experiments. *The annals of mathematical statistics*, 265–272.

Bolotnyy, V. and S. Vasserman (2023). Scaling auctions as insurance: A case study in infrastructure procurement. *Econometrica 91*(4), 1205–1259.

Che, Y.-K. (1993). Design competition through multidimensional auctions. *The RAND Journal of Economics*, 668–680.

Corner, J. L. and P. D. Corner (1995). Characteristics of decisions in decision analysis practice. *Journal of the Operational Research Society 46*(3), 304–314.

DeMarzo, P., I. Kremer, and A. Skrzypacz (2005a). Bidding with Securities: Auctions and Security Design. *American Economic Review 95*(4), 936–59.

DeMarzo, P. M., I. Kremer, and A. Skrzypacz (2005b). Bidding with securities: Auctions and security design. *American economic review 95*(4), 936–959.

Fioriti, A. and A. Hernandez-Chanto (2022). Leveling the playing field for risk-averse agents in security-bid auctions. *Management Science 68*(7), 5441–5463.

Fok, C. I. (2023). *Three Essays in Industrial Organization.* Phd thesis, University of Toronto (Canada).

Hanazono, M., J. Nakabayashi, R. Sano, and M. Tsuruoka (2024). Beyond quasilinearity: Exploring nonlinear scoring rules in procurement auctions. Working paper, University of Tokyo Market Design Center.

Hortaçsu, A. and S. L. Puller (2008). Understanding strategic bidding in multi-unit auctions: a case study of the texas electricity spot market. *The RAND Journal of Economics 39*(1), 86–114.

Krasnokutskaya, E. (2011). Identification and estimation of auction models with unobserved heterogeneity. *The Review of Economic Studies 78*(1), 293–327.

Krishna, V. (2009). *Auction theory.* Academic press.

Luo, Y. and H. Takahashi (2025). Bidding for contracts under uncertain demand: Skewed bidding and risk sharing. *Rand Journal of Economics 56*(3), 325–343.

Manelli, A. M. and D. R. Vincent (2006). Bundling as an optimal selling mechanism for a multiple-good monopolist. *Journal of Economic Theory 127*(1), 1–35.

Mas-Colell, A., M. D. Whinston, J. R. Green, et al. (1995). *Microeconomic theory*, Volume 1. Oxford university press New York.

Maskin, E. and J. Riley (1984). Optimal auctions with risk averse buyers. *Econometrica: Journal of the Econometric Society*, 1473–1518.

Matthews, S. (1987). Comparing auctions for risk averse buyers: A buyer's point of view. *Econometrica: Journal of the Econometric Society*, 633–646.

Mehra, R. and E. C. Prescott (1985). The equity premium: A puzzle. *Journal of Monetary Economics 15*, 145–161.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of operations research 6*(1), 58–73.

Riley, J. G. and W. F. Samuelson (1981). Optimal auctions. *The American Economic Review 71*(3), 381–392.

Topkis, D. M. (1978). Minimizing a submodular function on a lattice. *Operations Research 26*(2), 305–321.

Vives, X. (1999). *Oligopoly pricing: old ideas and new tools.* MIT press.

A Omitted Proofs

Proof of Lemma 1

The continuity of $U(\cdot)$ follows from the fact that $u(\cdot)$ is continuous. To show strict concavity, fix $\boldsymbol{b}_i,\boldsymbol{b}_i' \in \mathbb{R}^T$ with $\boldsymbol{b}_i \neq \boldsymbol{b}_i'$ and $\lambda \in (0,1)$. Notice that

$$\mu\Big(\Big\{\boldsymbol{q} : \boldsymbol{q}\cdot(\boldsymbol{b}_i - \boldsymbol{c}) - k_i = \boldsymbol{q}\cdot(\boldsymbol{b}_i' - \boldsymbol{c}) - k_i\Big\}\Big) = \mu\Big(\Big\{\boldsymbol{q} : \boldsymbol{q}\cdot(\boldsymbol{b}_i - \boldsymbol{b}_i') = 0\Big\}\Big) < 1,$$

where the inequality follows from the fact that μ is generic. Thus,

$$\begin{aligned}
U\big(\lambda\boldsymbol{b}_i + (1-\lambda)\boldsymbol{b}_i' \mid k_i\big) &:= \mathbb{E}_\mu\big[u\big(\boldsymbol{q}\cdot\big(\lambda\boldsymbol{b}_i + (1-\lambda)\boldsymbol{b}_i' - \boldsymbol{c}\big) - k_i\big)\big] \\
&= \mathbb{E}_\mu\big[u\big(\lambda\big(\boldsymbol{q}\cdot(\boldsymbol{b}_i - \boldsymbol{c}) - k_i\big) + (1-\lambda)\big(\boldsymbol{q}\cdot(\boldsymbol{b}_i' - \boldsymbol{c}) - k_i\big)\big)\big] \\
&> \lambda\cdot\mathbb{E}_\mu\big[u\big(\boldsymbol{q}\cdot(\boldsymbol{b}_i - \boldsymbol{c}) - k_i\big)\big] + (1-\lambda)\cdot\mathbb{E}_\mu\big[u\big(\boldsymbol{q}\cdot(\boldsymbol{b}_i' - \boldsymbol{c}) - k_i\big)\big] \\
&= \lambda\cdot U(\boldsymbol{b}_i \mid k_i) + (1-\lambda)\cdot U(\boldsymbol{b}_i' \mid k_i).
\end{aligned}$$

where the strict inequality follows from Jensen's inequality.

Now, assume the support of μ is bounded and contained in $\mathbb{R}^T_+$. Observe that if $b_{i,t}$ dominates $b'_{i,t}$ in the component-wise order, then for each $\boldsymbol{q} \in \mathbb{R}^T_+$, $u\left(\boldsymbol{q}\cdot(\boldsymbol{b}_i - \boldsymbol{c}) - k_i\right) > u\left(\boldsymbol{q}\cdot(\boldsymbol{b}_i' - \boldsymbol{c}) - k_i\right)$, which implies that $U\left(\boldsymbol{b}_i \mid k_i\right) \geqslant U\left(\boldsymbol{b}_i' \mid k_i\right)$.

We show that $U(\cdot \mid k_i)$ is differentiable. Recall u is continuously differentiable. Moreover, for each $\boldsymbol{b}_i \in \mathbb{R}^T_+$, since μ has bounded support, there is some $K \in \mathbb{R}$ such that $K < u\left(\boldsymbol{q}\cdot(\boldsymbol{b}_i - \boldsymbol{c}) - k_i\right) < 1$ almost surely. So, the dominated convergence theorem applies. Thus, $U(\cdot \mid k_i)$ is differentiable with partial derivatives given by

$$\frac{\partial}{\partial b_{i,t}}\mathbb{E}_\mu\left[u\left(\sum_{t=1}^{T} q_t\left(b_{i,t} - c_t\right) - k_i\right)\right] = \mathbb{E}_\mu\left[\frac{\partial}{\partial b_{i,t}} u\left(\sum_{t=1}^{T} q_t\left(b'_{i,t} - c_t\right) - k_i\right)\right].$$

The dominated convergence theorem implies that the derivative of $U(\cdot)$ is continuous.

Proof of Lemma 2

Fix a scaling auction $\mathcal{A} = (\boldsymbol{w}, \mathcal{B})$ with costs $k_i, k_i' \in [\underline{k}, \overline{k}]$, and a feasible score s_i. To show that $\boldsymbol{b}_i(s_i \mid k_i, \mathcal{A}) = \boldsymbol{b}_i(s_i \mid k_i', \mathcal{A})$, it suffices to show that k_i and k_i' share the same preferences over bidding vectors. Fix $\boldsymbol{b}_i, \boldsymbol{b}_i' \in \mathcal{B}$. Since u has no wealth effects, $\mathbb{E}_\mu\left[u\left(\boldsymbol{q} \cdot (\boldsymbol{b}_i - \boldsymbol{c}) - k_i\right)\right] \geqslant \mathbb{E}_\mu\left[u\left(\boldsymbol{q} \cdot (\boldsymbol{b}_i' - \boldsymbol{c}) - k_i\right)\right]$ if and only if $\mathbb{E}_\mu\left[u\left(\boldsymbol{q} \cdot (\boldsymbol{b}_i - \boldsymbol{c}) - k_i'\right)\right] \geqslant \mathbb{E}_\mu\left[u\left(\boldsymbol{q} \cdot (\boldsymbol{b}_i' - \boldsymbol{c}) - k_i'\right)\right]$. Thus, $U\left(\boldsymbol{b}_i \mid k_i\right) \geqslant U\left(\boldsymbol{b}_i' \mid k_i\right)$ if and only if $U\left(\boldsymbol{b}_i \mid k_i'\right) \geqslant U\left(\boldsymbol{b}_i' \mid k_i'\right)$.

Proof of Proposition 1

Lemma A.1. *Fix a rich scaling auction $\mathcal{A} = (\boldsymbol{w}, \mathcal{B})$. For each $k \in [\underline{k}, \overline{k}]$ there exists some score $s \in \mathbb{R}$ such that $\mathrm{CE}(s \mid \mathcal{A}) = \psi(k)$. Moreover, if the scores $s, s' \in \mathbb{R}$ are feasible for $\mathcal{A}$ and $s \leqslant s'$, then $\mathrm{CE}(s \mid \mathcal{A}) \leqslant \mathrm{CE}(s' \mid \mathcal{A})$.*

Proof. Fix $k \in [\underline{k}, \overline{k}]$. Observe that since ψ is increasing, $\psi(\overline{k}) = \overline{k}$, and $\psi(k) \geqslant k$, it follows that there is some $k' \in [\underline{k}, \overline{k}]$ such that $k' = \psi(k)$. Since $\mathcal{A}$ is rich, there is some feasible score $s \in \mathbb{R}$ such that $\mathrm{CE}(s \mid \mathcal{A}) = k' = \psi(k)$. Moreover, if $s \leqslant s'$, then the budget sets satisfy $\mathcal{B}[s \mid \boldsymbol{w}] \subseteq \mathcal{B}[s' \mid \boldsymbol{w}]$. Thus,

$$u\big(\mathrm{CE}(s \mid \mathcal{A})\big) = \mathbb{E}_\mu\big[u(\boldsymbol{q} \cdot (\boldsymbol{b}(s \mid \mathcal{A}) - \boldsymbol{c}))\big] \leqslant \mathbb{E}_\mu\big[u(\boldsymbol{q} \cdot (\boldsymbol{b}(s' \mid \mathcal{A}) - \boldsymbol{c}))\big] = u\big(\mathrm{CE}(s' \mid \mathcal{A})\big),$$

where the equalities follow from the definition of $\mathrm{CE}(\cdot \mid \mathcal{A})$, and the inequality follows from the fact that $\mathcal{B}[s \mid \boldsymbol{w}] \subseteq \mathcal{B}[s' \mid \boldsymbol{w}]$. Since $u(\cdot)$ is strictly increasing, it follows that $\mathrm{CE}(s \mid \mathcal{A}) \leqslant \mathrm{CE}(s' \mid \mathcal{A})$. $\square$

Lemma A.2. *Fix a rich scaling auction $\mathcal{A} = (\boldsymbol{w}, \mathcal{B})$ and let $\varphi_{\mathcal{A}} : [\underline{k}, \overline{k}] \to \mathbb{R}$ be an increasing mapping such that $\mathrm{CE}\big(\varphi_{\mathcal{A}}(k_i) \mid \mathcal{A}\big) = \psi(k_i)$. The following hold:*

1. *For each $k_i \in [\underline{k}, \overline{k}]$ and $s_i \in \mathbb{R}$, $\mathrm{EU}(s_i, \mathcal{A} \mid k_i) = \mathrm{EU}_{\mathrm{aux}}(\mathrm{CE}(s_i \mid \mathcal{A}) \mid k_i)$.*
2. *For each $k_i \in [\underline{k}, \overline{k}]$, $\mathrm{EU}(\varphi_{\mathcal{A}}(k_i), \mathcal{A} \mid k_i) = \mathrm{EU}_{\mathrm{aux}}(\psi(k_i) \mid k_i)$.*

Proof. Write $S \equiv \varphi_{\mathcal{A}}\big([\underline{k}, \overline{k}]\big)$ for the range of $\varphi_{\mathcal{A}}$ and let $s_i \in S$. Since $\psi(\cdot)$ and $\mathrm{CE}(\cdot \mid \mathcal{A})$ are increasing, then $\psi^{-1}(\mathrm{CE}(s_i \mid \mathcal{A})) = \varphi_{\mathcal{A}}^{-1}(s_i)$. Consequently, for each

$j \neq i$, $\Pr[s_i \leqslant \varphi_{\mathcal{A}}(k_j)] = \Pr\big[\mathrm{CE}(s_i \mid \mathcal{A}) \leqslant \psi(k_j)\big]$. Hence, for each $k_i \in [\underline{k}, \overline{k}]$,

$$\begin{aligned}\mathrm{EU}_{\mathrm{aux}}\big(\mathrm{CE}(s_i \mid \mathcal{A}) \mid k_i\big) &= u\big(\mathrm{CE}(s_i \mid \mathcal{A}) - k_i\big) \cdot \prod_{j \neq i} \Pr\big[\mathrm{CE}(s_i \mid \mathcal{A}) \leqslant \psi(k_j)\big] \\ &= U(s_i \mid k_i, \mathcal{A}) \cdot \prod_{j \neq i} \Pr\big[s_i \leqslant \varphi_{\mathcal{A}}(k_j)\big] \\ &= \mathrm{EU}(s_i, \mathcal{A} \mid k_i).\end{aligned}$$

Taking $s_i = \varphi_{\mathcal{A}}(k_i)$ implies $\mathrm{EU}(\varphi_{\mathcal{A}}(k_i) \mid k_i, \mathcal{A}) = \mathrm{EU}_{\mathrm{aux}}(\psi(k_i) \mid k)$. $\square$

Fix a rich scaling auction $\mathcal{A} = (\boldsymbol{w}, \mathcal{B})$. We first show the existence of a symmetric equilibrium $\varphi_{\mathcal{A}}$ for $\mathcal{A}$. Write $\mathcal{S}$ for the set of feasible scores for $\mathcal{A}$. Write $S(\psi(k_i)) := \{s_i \in \mathcal{S} : \mathrm{CE}(s_i \mid \mathcal{A}) = \psi(k_i)\}$. Notice, since $\mathcal{A}$ is rich, $S(\psi(k_i))$ is non empty for each $k_i \in [\underline{k}, \overline{k}]$ (See Lemma A.1). Define the strategy $\varphi_{\mathcal{A}} : [\underline{k}, \overline{k}] \to \mathbb{R}$ given by $\varphi_{\mathcal{A}}(k_i) = \min\{S(\psi(k_i))\}$. So, $\mathrm{CE}\big(\varphi_{\mathcal{A}}(k_i) \mid \mathcal{A}\big) = \psi(k_i)$. Observe that $\mathrm{CE}(\cdot \mid \mathcal{A})$ is increasing (See Lemma A.1). Since $\psi(\cdot)$ is strictly increasing, then $\varphi_{\mathcal{A}}(\cdot)$ is strictly increasing. Fix $k_i \in [\underline{k}, \overline{k}]$ and a score $s_i \in \mathbb{R}$ that is feasible for $\mathcal{A}$. Observe that

$$\mathrm{EU}\big(\varphi_{\mathcal{A}}(k_i) \mid k_i, \mathcal{A}\big) = \mathrm{EU}_{\mathrm{aux}}\big(\psi(k_i) \mid k_i\big) \geqslant \mathrm{EU}_{\mathrm{aux}}\big(\mathrm{CE}(s_i \mid \mathcal{A}) \mid k_i\big) = \mathrm{EU}(s_i \mid k_i, \mathcal{A}),$$

where the first and third equalities follow from Lemma A.2 and the inequality follows from the fact that $\psi(\cdot)$ is an equilibrium of the auxiliary auction. Therefore, $\varphi_{\mathcal{A}}$ is an equilibrium of the scaling auction $\mathcal{A}$.

Now, we show uniqueness. Let $\hat{\varphi}_{\mathcal{A}} : [\underline{k}, \overline{k}] \to \mathbb{R}$ be an equilibrium of $\mathcal{A}$. Define $\hat{\psi} : [\underline{k}, \overline{k}] \to \mathbb{R}$ as $\hat{\psi}(k_i) := \mathrm{CE}(\hat{\varphi}_{\mathcal{A}}(k_i) \mid \mathcal{A})$. Fix r_i in the range of $\mathrm{CE}(\cdot \mid \mathcal{A})$ and let $s_i \in \mathbb{R}$ such that $r_i = \mathrm{CE}(s_i \mid \mathcal{A})$. Thus,

$$\mathrm{EU}_{\mathrm{aux}}\big(\hat{\psi}(k_i) \mid k_i\big) = \mathrm{EU}\big(\hat{\varphi}_{\mathcal{A}}(k_i) \mid k_i, \mathcal{A}\big) \geqslant \mathrm{EU}(s_i \mid k_i, \mathcal{A}) = \mathrm{EU}_{\mathrm{aux}}\big(\mathrm{CE}(s_i \mid \mathcal{A}) \mid k_i\big) = \mathrm{EU}_{\mathrm{aux}}(r_i \mid k_i)$$

where the first and second equalities follow from Lemma A.2 and the inequality holds because $\hat{\varphi}_{\mathcal{A}}(\cdot)$ is an equilibrium. Hence, $\hat{\psi}(\cdot)$ is an equilibrium of the auxiliary auction. However, since the auxiliary auction has a unique symmetric equilibrium, then $\hat{\psi} = \psi$. This implies that $\mathrm{CE}(\hat{\varphi}_{\mathcal{A}}(k_i) \mid \mathcal{A}) = \mathrm{CE}(\varphi_{\mathcal{A}}(k_i) \mid \mathcal{A})$ for each $k_i \in [\underline{k}, \overline{k}]$. Consequently, $\hat{\varphi}_{\mathcal{A}}(k_i) = \varphi_{\mathcal{A}}(k_i)$ for each $k_i \in [\underline{k}, \overline{k}]$, as desired.

Proof of Theorem 1

Fix a rich scaling auction $\mathcal{A} = (\boldsymbol{w}, \mathcal{B})$ such that $\mathcal{B} = \mathbb{R}^T$ and $\boldsymbol{w} = \lambda \cdot \mathbb{E}_\mu[\boldsymbol{q}]$ for some $\lambda > 0$. Let $k \in [\underline{k}, \overline{k}]$. Define $\boldsymbol{b}_i^* = \boldsymbol{b}_i(\varphi_{\mathcal{A}}(k_i) \mid \mathcal{A})$ and $\bar{u}_i = u(\psi(k_i) - k_i)$. Consider another rich scaling auction $\mathcal{A}' = (\boldsymbol{w}', \mathcal{B}')$ and let $\boldsymbol{b}_i' = \boldsymbol{b}_i(\varphi_{\mathcal{A}'}(k_i) \mid \mathcal{A}')$. We will show that $\mathbb{E}_\mu[\boldsymbol{b}_i^* \cdot \boldsymbol{q}] \leqslant \mathbb{E}_\mu[\boldsymbol{b}_i' \cdot \boldsymbol{q}]$. Note, by Proposition 1, we have $\bar{u}_i = U(\boldsymbol{b}_i^* \mid k_i) = U(\boldsymbol{b}_i' \mid k_i)$. Since $\mathcal{B}' \subseteq \mathbb{R}^T = \mathcal{B}$, it follows that $\boldsymbol{b}_i' \in B(\bar{u}_i \mid k_i, \mathcal{A})$. Hence,

$$\lambda\, \mathbb{E}_\mu[\boldsymbol{b}_i^* \cdot \boldsymbol{q}] = \boldsymbol{b}_i^* \cdot \boldsymbol{w} \leqslant \boldsymbol{b}_i' \cdot \boldsymbol{w} = \lambda\, \mathbb{E}_\mu[\boldsymbol{b}_i' \cdot \boldsymbol{q}],$$

where the equality follows from the fact that $\boldsymbol{w} = \lambda \mathbb{E}_\mu[\boldsymbol{q}]$ and the inequality follows from Proposition 2. Therefore, $\mathcal{A}$ is optimal.

Proof of Lemma 4

Fix $r, r' \in \mathbb{R}$, $k_i \in \left[\underline{k}, \overline{k}\right]$, and observe that $u(r - k_i - z) = u(r' - k_i)$ if and only if $r = r' + z$. So, the game of the auxiliary auction is equivalent to the game of the cash auction by shifting the actions via the mapping $g(r) \equiv r + z$. Consequently, $\zeta : [\underline{k}, \overline{k}] \to \mathbb{R}$ is an equilibrium of the scaling auction if and only if $\zeta(k_i) = \psi(k_i) + z$.

Proof of Proposition 3

Fix a cost $k_i \in [\underline{k}, \overline{k}]$ and note that, $u\big(\zeta(k_i) - k_i - z\big) = u(\psi(k_i) - k_i)$. (See Lemma 4.) As a consequence, each seller with fixed cost k_i expects the same utility in both auctions conditional on winning. Moreover, since ψ and ζ are both strictly increasing, the probability of winning is the same in both auctions. Thus, all seller' types expect the same utility in the auxiliary and the cash auction. So, by applying Corollary 1, sellers also expect the same expected utility as in each scaling auction.

Proof of Lemma 5

Fix $k_i \in [\underline{k}, \overline{k}]$, write $\overline{u}_i = u\big(\psi(k_i) - k_i\big)$. The set of markups that lead to utility $\overline{u}_i$ given fixed cost k_i is $\mathcal{M}(\overline{u}_i \mid k_i) \equiv \{\hat{\boldsymbol{m}}_i \in \mathbb{R}^T : \mathbb{E}_\mu\left[u\left(\boldsymbol{q} \cdot \hat{\boldsymbol{m}}_i - k_i\right)\right] = \overline{u}_i\}$. Notice, $B(\overline{u}_i \mid k_i, \mathcal{A}) = \mathcal{M}(\overline{u}_i \mid k_i) + \boldsymbol{c} \equiv \big\{\hat{\boldsymbol{m}}_i + \boldsymbol{c} : \hat{\boldsymbol{m}}_i \in \mathcal{M}(\overline{u}_i \mid k_i)\big\}$, i.e., $B(\overline{u}_i \mid k_i, \mathcal{A})$ is $\mathcal{M}(\overline{u}_i \mid k_i)$ translated by the vector $\boldsymbol{c}$.

Proposition 1 implies that $\boldsymbol{b}_i(\varphi_{\mathcal{A}}(k_i) \mid \mathcal{A}) \in B(\overline{u}_i \mid k_i, \mathcal{A})$. If $\mathcal{A}$ is optimal, then

$$\boldsymbol{b}_i(\varphi_{\mathcal{A}}(k_i) \mid \mathcal{A}) \in \arg \min_{\hat{\boldsymbol{b}}_i \in B(\overline{u}_i \mid k_i, \mathcal{A})} \mathbb{E}_\mu\big[\hat{\boldsymbol{b}}_i \cdot \boldsymbol{q}\big].$$

(See Theorem 1.) Moreover, Since, $\boldsymbol{b}_i(\varphi_{\mathcal{A}}(k_i) \mid \mathcal{A}) = \boldsymbol{m}_i(k_i \mid \mathcal{A}) + \boldsymbol{c}$ and $B(\overline{u}_i \mid k_i, \mathcal{A}) = \mathcal{M}(\overline{u}_i \mid k_i) + \boldsymbol{c}$, the minimization problem above is equivalent to require

$$\boldsymbol{m}_i(k_i \mid \mathcal{A}) \in \arg \min_{\hat{\boldsymbol{m}}_i \in \mathcal{M}(\overline{u}_i \mid k_i)} \mathbb{E}_\mu\big[\hat{\boldsymbol{m}}_i \cdot \boldsymbol{q}\big].$$

Proof of Theorem 2

Lemma A.3. *Assume the buyer implements a rich scaling auction $\mathcal{A}$. If $k_i \in [\underline{k}, \overline{k}]$ is the fixed cost of the winning seller, then the buyer's expected payment is*

$$\mathbb{E}_\mu[\boldsymbol{q} \cdot \boldsymbol{b}_i(\varphi_{\mathcal{A}}(k_i) \mid \mathcal{A})] = \psi(k_i) + \mathbb{E}_\mu[\boldsymbol{q} \cdot \boldsymbol{c}] + \mathrm{RP}(\boldsymbol{m}_i(k_i \mid \mathcal{A})).$$

Proof. Fix $k_i \in [\underline{k}, \overline{k}]$ and write $\boldsymbol{b}_i^* = \boldsymbol{b}_i(\varphi_{\mathcal{A}}(k_i) \mid \mathcal{A})$. By Proposition 3,

$$\mathbb{E}_\mu\big[u\big(\boldsymbol{q} \cdot (\boldsymbol{b}_i^* - \boldsymbol{c} - k_i)\big)\big] = U(\boldsymbol{b}_i^* \mid k_i) = u\big(\psi(k_i) - k_i\big).$$

Moreover, since $u(\cdot)$ has no wealth effects, $\mathbb{E}_\mu\big[u\big(\boldsymbol{q} \cdot (\boldsymbol{b}_i^* - \boldsymbol{c})\big)\big] = u\big(\psi(k_i)\big)$. Therefore, $\mathbb{E}_\mu[\boldsymbol{q} \cdot (\boldsymbol{b}_i^* - \boldsymbol{c})] - \mathrm{RP}(\boldsymbol{b}_i^* - \boldsymbol{c}) = \psi(k_i)$, which implies the desired equality. □

Lemma A.4. *Assume the buyer implements a cash auction. If $k_i \in [\underline{k}, \overline{k}]$ is the type of the winning seller, then the buyer's payment is $\zeta(k_i) = \psi(k_i) + \mathbb{E}_\mu[\boldsymbol{q} \cdot \boldsymbol{c}] + \mathrm{RP}(-\boldsymbol{c})$.*

Proof. The proof follows from Lemma 4 and the fact that $z = \mathbb{E}_\mu[\boldsymbol{q} \cdot \boldsymbol{c}] + \mathrm{RP}(-\boldsymbol{c})$. □

The first part follows directly from Lemmata A.3 and A.4. We now show the second part. Write $\overline{R} \equiv \sup_{k_i \in [\underline{k}, \overline{k}]} \mathrm{RP}(\boldsymbol{m}_i(k_i \mid \mathcal{A}))$, for the maximum risk faced by any seller. Likewise, write $\underline{R} \equiv \inf_{k_i \in [\underline{k}, \overline{k}]} \mathrm{RP}(\boldsymbol{m}_i(k_i \mid \mathcal{A}))$. Note, the values $0 \leqslant \underline{R} \leqslant \overline{R}$ are independent of $\boldsymbol{c}$. (See Lemma 5). Now, notice that $\mathrm{RP}(\cdot)$ is continuous and that $\mathrm{RP}(\mathbf{0}) = 0$. Thus, if the marginal costs $\boldsymbol{c}$ are sufficiently close to zero, then $\mathrm{RP}(-\boldsymbol{c}) \leqslant \underline{R}$. Conversely, if the marginal costs $\boldsymbol{c}$ are large enough, then $\mathrm{RP}(-\boldsymbol{c}) \geqslant \overline{R}$.

Proof of Theorem 3

Lemma A.5. *Fix a hybrid scaling auction* $\hat{\mathcal{A}} = (\hat{\boldsymbol{w}}, \hat{\mathcal{B}})$ *in which* $\hat{\boldsymbol{w}} = \big(1, \mathbb{E}_\mu[q_1], \ldots, \mathbb{E}_\mu[q_T]\big)$ *and* $\mathbb{R}^T_+ \subseteq \hat{\mathcal{B}}$. *Write* $\rho(s) = s - \mathbb{E}_{\hat{\mu}}[\hat{\boldsymbol{q}}] \cdot \hat{\boldsymbol{c}}$. *If* $\rho(s) \geqslant 0$, *then* $\hat{\boldsymbol{b}}(s \mid \hat{\mathcal{A}}) = (\rho(s), c_1, \ldots, c_T)$ *and* $\mathrm{CE}(s \mid \hat{\mathcal{A}}) = \rho(s)$.

Proof. Fix $k \in [\underline{k}, \overline{k}]$. Observe that the utility function $U(\cdot \mid k)$ is strictly concave, strictly monotone, and differentiable, due to Lemma 1. Consequently, the first-order condition is sufficient to characterize the optimal bidding vector. Since there are no wealth effects, $u(\pi) := 1 - \exp(-\gamma\pi)$ for some $\gamma > 0$. So, the Lagrangian is $\mathcal{L}(\hat{\boldsymbol{b}}, \lambda) \equiv U(\hat{\boldsymbol{b}} \mid k) - \lambda\left(\hat{\boldsymbol{b}} \cdot \mathbb{E}_{\hat{\mu}}[\hat{\boldsymbol{q}}] - s\right)$. Then, for each $t \in \mathcal{T}$,

$$\tfrac{\partial \mathcal{L}(\hat{\boldsymbol{b}}, \lambda)}{\partial b_t} = \mathbb{E}_{\hat{\mu}}\left[\gamma \exp\left(-\gamma(\hat{\boldsymbol{b}} \cdot \hat{\boldsymbol{q}} - \hat{\boldsymbol{c}} \cdot \hat{\boldsymbol{q}} - k)\right) q_t\right] - \lambda \mathbb{E}_{\hat{\mu}}[q_t].$$

Let $\hat{\boldsymbol{b}}^* = (\rho(s), c_1, \ldots, c_T)$ and $\lambda^* = \gamma \exp\left(-\gamma(\rho(s) - k)\right)$ and notice that $\hat{\boldsymbol{b}}^* \cdot \hat{\boldsymbol{w}} = s$. Moreover, since $q_0 = 1$ almost surely, we have $\hat{\boldsymbol{b}}^* \cdot \hat{\boldsymbol{q}} - \hat{\boldsymbol{c}} \cdot \hat{\boldsymbol{q}} = \rho(s)$. Then, if $t \geqslant 1$,

$$\tfrac{\partial \mathcal{L}(\hat{\boldsymbol{b}}^*, \lambda^*)}{\partial b_t} = \gamma \exp(-\gamma(\rho(s) - k)) \cdot \mathbb{E}_{\hat{\mu}}[q_t] - \lambda^* \mathbb{E}_{\hat{\mu}}[q_t] = 0.$$

Similarly, if $t = 0$, then $\frac{\partial \mathcal{L}(\hat{\boldsymbol{b}}^*, \lambda^*)}{\partial b_0} = \gamma \exp(-\gamma(\rho(s) - k)) - \lambda^* = 0$. Thus, $(\hat{\boldsymbol{b}}^*, \lambda^*)$ satisfies the first-order conditions, and $\hat{\boldsymbol{b}}^*$ maximizes $U(\cdot \mid k)$ subject to $\hat{\boldsymbol{b}} \cdot \hat{\boldsymbol{w}} = s$. Since $\rho(s) \geqslant 0$, it follows that $\hat{\boldsymbol{b}}^* \in \mathbb{R}^T_+ \subseteq \hat{\mathcal{B}}$. Thus, $\mathrm{CE}(s \mid \hat{\mathcal{A}}) = \rho(s)$. □

First, we show that $\hat{\mathcal{A}}$ is rich and optimal. Consider any fixed $k_i \in [\underline{k}, \overline{k}]$. If $s_i = k_i + \mathbb{E}_{\hat{\mu}}[\hat{\boldsymbol{q}}] \cdot \hat{\boldsymbol{c}}$, then $\mathrm{CE}(s_i \mid \hat{\mathcal{A}}) = k_i$ (see Lemma A.5). This shows that $\hat{\mathcal{A}}$

is rich. Second, we show that $\hat{\mathcal{A}}$ provides full insurance. By Lemma A.5, we have $\hat{\boldsymbol{m}}(k_i \mid \hat{\mathcal{A}}) \cdot \hat{\boldsymbol{q}} = \rho(s_i)$ almost surely. Consequently, $\mathrm{RP}(\hat{\boldsymbol{m}}(k_i \mid \hat{\mathcal{A}}) \cdot \hat{\boldsymbol{q}}) = 0$. Third, we show that $\hat{\mathcal{A}}$ induces a lower payment than any other scaling auction (hybrid or not) or the cash auction. Let k_i be the fixed cost of the winning seller. Observe that the buyer's expected payment is $\psi(k_i) + \mathbb{E}_{\hat{\mu}}[\hat{\boldsymbol{q}} \cdot \boldsymbol{c}]$. (See Lemma A.3). This expected payment is lower than the payment made by any other scaling auction (hybrid or standard) and the expected payment made by the cash auction (see Lemmata A.3 and A.4).

Proof of Lemma 6

Let $\psi : [\underline{k}, \kappa_{\mathcal{A}}] \to \mathbb{R}$ be a mapping such that $\mathrm{CE}\big(\varphi_{\mathcal{A}}(k_i) \mid \mathcal{A}\big) = \psi(k_i)$. We show that ψ is an equilibrium of the auxiliary auction under $R_{\mathcal{A}}$. Lemma A.2 holds in this setting as well. So, the result holds using an analogous argument as Proposition 1.

Proof of Proposition 4

By Lemma 1, we know that the preferences induced by $U(\cdot \mid \theta)$ are strictly convex and strictly monotone. Define $B_0 \equiv \bigcup_\theta B_0(\theta)$ as the set of bids that result in zero utility for some type of seller. write $B^* = \arg\min_{B_0} \boldsymbol{b} \cdot \mathbb{E}_\mu[\boldsymbol{q}]$ for the set of bids that minimize the score and leave sellers indifferent to participating in the auction.

Fix $\boldsymbol{b}^* \in B^*$ and let $\theta \in \Theta$ be a type such that $\boldsymbol{b}^* \in B_0(\theta)$. Consider a strategy profile where type θ bids $\boldsymbol{b}^*$ and all other sellers choose bids with higher scores. Note that beating a bid of $\boldsymbol{b}^*$ would require a bid that results in negative utility for all types. Hence, no type has an incentive to deviate from this strategy, ensuring that this profile constitutes a Bayesian equilibrium. Next, we show that no other scaling auction can reduce the expected payment further. Note, for each type θ^*, $U(\boldsymbol{b}' \mid \theta^*) \leqslant 0$. Therefore, if a scaling auction induces a bid $\boldsymbol{b}'$ such that $\boldsymbol{b}' \cdot \mathbb{E}_\mu[\boldsymbol{q}] < \boldsymbol{b}^* \cdot \mathbb{E}_\mu[\boldsymbol{q}]$, then by strict monotonicity, $U(\boldsymbol{b}' \mid \theta) < 0$ for some type θ. Thus, no seller would have an incentive to choose such a bid, as it would lead to negative utility.

Supplementary Appendix B

For Online Publication Only

This Appendix analyzes a broad class of scoring auctions for procurement in settings with risk-averse sellers who possess private information about their costs and input requirements. We compare various scoring auction formats (beyond scaling auctions) and identify conditions under which sellers' utilities are equivalent across these formats. Our main result establishes that two individually rational scoring auctions deliver identical expected utility to every seller type if and only if they satisfy a single-crossing condition, which we call *homogeneous type effects*. We also present two sufficient conditions—*separability* and *competitiveness*—that are easier to verify and encompass the settings discussed in the main text.

B.1 Environment

A buyer requires the completion of a project. To this end, the buyer organizes a procurement auction. In this competitive setting, a finite set of qualified sellers $\mathcal{I} = \{1, 2, \ldots, I\}$ compete for the contract to execute the project. (Section B.4.2 extends the analysis to the case in which $\mathcal{I}$ is infinite.)

The project consists of a set of tasks $\mathcal{T} = \{1, 2, \ldots, T\}$ defined by the buyer at the outset of the bidding process, all of which must be completed for the project's successful execution. Each task $t \in \mathcal{T}$ requires different input quantities, denoted by the vector $\boldsymbol{q} = (q_t : t \in \mathcal{T})$, which are stochastic. Importantly, the realization of $\boldsymbol{q}$ is independent of the auction outcome—such as the identity of the winning bidder or the final payment.

Prior to the start of the project, each seller $i \in \mathcal{I}$ privately learns their efficiency type $\theta_i \in \Theta := [\underline{\theta}, \overline{\theta}]$, drawn independently from a commonly known, atomless distribution F with density f. A seller's type θ_i governs three elements of their cost structure and information. First, it determines their fixed cost $k(\theta_i)$, which captures organizational or managerial inefficiencies internal to the seller and unrelated to input prices. Second, it shapes the seller's marginal costs $c_t(\theta_i)$ of procuring the quantity q_t for each task t, reflecting type-specific efficiencies or bargaining power in input markets. The vector of marginal costs for seller i is denoted $\boldsymbol{c}(\theta_i) = (c_t(\theta_i) : t \in \mathcal{T})$, and these costs are private to the seller. Third, the seller's type determines their distribution of input requirements: each seller has a distribution μ over the realization of the input vector $\boldsymbol{q}$, with $\mu(\theta_i) \in \Delta(\mathbb{R}^T_+)$ denoting the conditional distribution given type θ_i.

Sellers evaluate their final monetary outcome through a Bernoulli utility function $u : \mathbb{R} \to \mathbb{R}$, applied to auction revenue net of project delivery costs. The utility function u is strictly increasing, weakly concave, and normalized so that $u(0) = 0$. By contrast, the buyer is risk-neutral and seeks to minimize the expected payment.

An environment is the tuple $\mathcal{E} = \big(\mathcal{I}, [\underline{\theta}, \overline{\theta}], k(\cdot), c(\cdot), \mu(\cdot), u(\cdot)\big)$, which specifies the number of bidders, their project costs, their risk preferences, and their common distribution μ. Throughout the paper, this environment is taken as fixed.

B.2 Scoring auctions

Scoring auctions are mechanisms that allow the buyer to design a payment scheme contingent on the project's realization. Sellers compete by submitting bids in the form of scores, and the project is awarded to the bidder proposing the lowest score. The winning score then determines the final payment to the auction winner as a function of the observable vector of tasks $\boldsymbol{q}$.

A scoring auction, denoted $\mathcal{A}$, is characterized by three components: an interval of permissible bidding scores $[\underline{s}, \overline{s}]$; a compact-valued correspondence of alternatives $(B(s))_{s \in [\underline{s}, \overline{s}]}$; and a continuous payment rule $\Pi(\boldsymbol{q}, s, b)$ that specifies the buyer's contingent payment to the seller as a function of the realized quantity vector $\boldsymbol{q}$, the score s, and the alternative $b \in B(s)$.

The timing of a scoring auction unfolds as follows. First, all sellers simultaneously submit their bids: each seller i chooses a score s_i and an alternative $b_i \in B(s_i)$. The project is then awarded to the bidder with the lowest score.[34] Next, the winning seller i, of type θ_i, undertakes the project, incurring a private fixed cost $k(\theta_i)$ and a variable cost $\boldsymbol{q} \cdot \boldsymbol{c}(\theta_i)$. Finally, the buyer compensates the winner with a total payment $\Pi(\boldsymbol{q}, s_i, b_i)$. The winner's total ex-post monetary payoff is therefore

$$\pi_i = \Pi(\boldsymbol{q}, s_i, b_i) - \boldsymbol{q} \cdot \boldsymbol{c}(\theta_i) - k(\theta_i).$$

This setup implicitly assumes verifiability of the input vector $\boldsymbol{q}$ required by the payment rule Π. We now turn to three important classes of scoring auctions.

Cash Auctions In a cash auction, bidders compete over a fixed payment s, independent of the realized quantities $\boldsymbol{q}$. This corresponds to the auction format with payment rule $\Pi(\boldsymbol{q}, s, b) = s$. For each score s, the set of alternatives $B(s)$ contains only a single trivial element. Hence, the winning bidder i receives a total payment of s_i, and her ex-post monetary payoff is

$$\pi_i = s_i - \boldsymbol{q} \cdot \boldsymbol{c}(\theta_i) - k(\theta_i).$$

Scaling auctions A *scaling auction* is a specific type of scoring auction in which buyers bid with individual prices for each task. A scaling auction is formally defined by a pair $(\boldsymbol{w}, \mathcal{B})$, in which $\boldsymbol{w} = (w_t : t \in \mathcal{T}) \in \mathbb{R}_+^T$ is a non-zero vector specifying the task-specific weights for the scoring rule, and $\mathcal{B} \subseteq \mathbb{R}^T$ constitutes the convex and closed set of admissible bidding vectors. The buyer determines these rules at the auction's outset and announces them upfront.

In a scaling auction, each seller i must submit a per-unit bid $b_{i,t}$ for every task t. These bids $\boldsymbol{b}_i \in \mathcal{B}$ are submitted simultaneously and confidentially. Each bidding vector $\boldsymbol{b}_i$ induces a score $s_i = \boldsymbol{b}_i \cdot \boldsymbol{w}$. The contract is awarded to the seller proposing the *lowest score*, who then executes the project.

To see that this auction fits within a scoring auction framework, note that each

[34] In case of a tie, the winner is randomly selected.

seller i submits a bid vector $\boldsymbol{b}_i \in \mathcal{B}$, which serves as their chosen "alternative." The score $s_i = \boldsymbol{b}_i \cdot \boldsymbol{w}$ is derived from this bid, with permissible scores typically forming an interval $[\underline{s}, \overline{s}]$. For a given score s, the available alternatives are $B(s) = \{\boldsymbol{b} \in \mathcal{B} \mid \boldsymbol{b} \cdot \boldsymbol{w} = s\}$. The winning seller i receives a total payment $\Pi(\boldsymbol{q}, s_i, \boldsymbol{b}_i) = \boldsymbol{q} \cdot \boldsymbol{b}_i$, where $\boldsymbol{q} \cdot \boldsymbol{b}_i$ is the contingent payment.

Upon project completion, the winning seller i is compensated based on her unit bids $\boldsymbol{b}_i$ and the realized quantities $\boldsymbol{q}$. The total payment received from the buyer is $\Pi(\boldsymbol{q}, s_i, \boldsymbol{b}_i) = \boldsymbol{q} \cdot \boldsymbol{b}_i$.

Security-bid auctions A second prominent class of scoring auctions involves *security bids*. In this mechanism, the buyer specifies a *family of securities*, possibly complemented by a fixed upfront transfer $X \in \mathbb{R}$. This family represents a collection of potential contingent payment rules, where the score s identifies the particular security provided to the winner.

In these auctions, the choice of a score uniquely determines the alternative. Thus, for each score s, the set of available alternatives $B(s)$ is a trivial singleton. The project is awarded to the seller proposing the lowest score. The winning seller i, having bid s_i, then receives the upfront payment X plus the contingent payment specified by the chosen security, evaluated at the realized quantities $\boldsymbol{q}$ and a pre-specified cost vector $\boldsymbol{c}_*$ estimated by the buyer.[35]

Common examples of security families include:

> **Debt:** The seller receives a fixed payment $s \in \mathbb{R}_+$ if the project's variable cost exceeds this amount, with the buyer covering any remaining costs. Otherwise, the buyer pays the full variable cost. The buyer's total payment is
>
> $$\Pi(\boldsymbol{q}, s) = X + \min\{s, \boldsymbol{q} \cdot \boldsymbol{c}_*\}.$$

[35] In settings where inputs are derived from a perfectly competitive market, the cost vector $\boldsymbol{c}$ is identical across sellers and independent of type. In such environments, it is natural to set $\boldsymbol{c}_* = \boldsymbol{c}$.

Equity: The seller is reimbursed for a fraction $s \geqslant 0$ of the estimated variable costs. The buyer's payment is

$$\Pi(\boldsymbol{q}, s) = X + s \cdot (\boldsymbol{q} \cdot \boldsymbol{c}_*).$$

Fixed equity: The buyer announces a fixed equity proportion λ upfront, and sellers compete by bidding a cash amount s. The project is awarded to the seller submitting the lowest cash bid. Upon completion, the buyer pays the winning seller her cash bid s plus the fixed proportion λ of realized variable costs $\boldsymbol{q} \cdot \boldsymbol{c}_*$. The total payment is

$$\Pi(\boldsymbol{q}, s) = s + \lambda \cdot (\boldsymbol{q} \cdot \boldsymbol{c}_*).$$

Call options: The seller is protected against high variable costs. If the realized cost $\boldsymbol{q} \cdot \boldsymbol{c}_*$ exceeds a predetermined strike price $\tilde{s} \geqslant 0$, the buyer covers the excess; if the cost is below $\tilde{s}$, the buyer makes no contingent payment related to this option.[36] The score is defined as $s = -\tilde{s}$, the negative of the strike price, so the winning firm is the one offering the lowest score (highest strike price). The total payment is

$$\Pi(\boldsymbol{q}, s) = X + \max\{0, \boldsymbol{q} \cdot \boldsymbol{c}_* - \tilde{s}\} = X + \max\{0, \boldsymbol{q} \cdot \boldsymbol{c}_* + s\}.$$

B.3 Equilibrium analysis

Fix a scoring auction $\mathcal{A}$. The environment $\mathcal{E}$ and the scoring auction $\mathcal{A}$ together induce a Bayesian game in which each firm observes only its own type. We are interested in analyzing the equilibrium behavior of this Bayesian game.

Let $[\underline{s}, \overline{s}]$ denote the scoring set of auction $\mathcal{A}$. For each score $s \in [\underline{s}, \overline{s}]$, the seller faces a lottery over total monetary payoffs, contingent on the realization of the input

[36] This structure is equivalent to an insurance contract with deductible $\tilde{s}$.

vector $\boldsymbol{q}$. Accordingly, define

$$U(s,\theta \mid \mathcal{A}) := \max_{b\in B(s)} \mathop{\mathbb{E}}_{\boldsymbol{q}\sim\mu(\theta)} \Big[u\big(\Pi(\boldsymbol{q},s,b) - \boldsymbol{q}\cdot\boldsymbol{c}(\theta) - k(\theta)\big)\Big], \tag{B.1}$$

as the seller's maximum expected utility conditional on winning, over all possible alternatives $b \in B(s)$, including production costs.[37]

We say that auction $\mathcal{A}$ is **rich** if, for each type θ, there exists some score $s \in S$ such that $U(s,\theta \mid \mathcal{A}) = 0$. In a rich auction, the range of available scores S is broad enough that: (1) participation is individually rational for all types of firms, and (2) under a sufficiently large number of competitors, all surplus can be extracted from all types.[38]

Note that equilibrium behavior of any type expecting positive utility is contained in

$$U_+ := \big\{(\theta,s) : U(s,\theta \mid \mathcal{A}) > 0\big\}.$$

Definition B.1. *Fix a rich auction $\mathcal{A}$. Say that $\mathcal{A}$ is **well-behaved** if the restriction of $U(s,\theta \mid \mathcal{A})$ to U_+ is strictly increasing in s, decreasing in θ, differentiable, and log-supermodular.*

Thus, in a well-behaved auction, more efficient sellers (those with lower type θ) obtain higher expected utility from submitting lower bids. This structure guarantees that equilibrium bidding strategies are increasing in type, so more efficient sellers bid more aggressively.

Consider now a symmetric, strictly increasing (and therefore efficient) bidding strategy $\varphi : \Theta \to S$ employed by the sellers. We analyze the strategic behavior of a seller i of type θ_i. Let s_i be the score submitted by seller i. Seller i wins if her score s_i is lower than all $I-1$ rivals' scores. Since the types θ_j for $j \neq i$ are drawn

[37] Since $B(s)$ is compact and Π is continuous, the maximum is attained.

[38] Tools such as reserve prices may reduce the government's expected payment, but at the cost of a positive probability that no firm procures the project. By contrast, rich auctions allocate the project with probability one.

independently from F, and φ is strictly increasing, the score $\varphi(\theta_j)$ of opponent j is less than or equal to s_i if and only if $\theta_j \leqslant \varphi^{-1}(s_i)$.

The probability that a single opponent j submits a score greater than s_i is $1 - F(\varphi^{-1}(s_i))$. Hence, the probability that seller i wins with score s_i is

$$\mathbb{P}(i \text{ wins} \mid s_i; \varphi) = \big(1 - F(\varphi^{-1}(s_i))\big)^{I-1}.$$

If seller i wins, she achieves expected utility $U(s_i, \theta_i \mid \mathcal{A})$; if she loses, her payoff is zero, yielding $u(0) = 0$. Thus, the expected utility for seller i of type θ_i submitting score s_i, when others follow strategy φ, is

$$\mathbb{P}(i \text{ wins} \mid s_i; \varphi) \cdot U(s_i, \theta_i \mid \mathcal{A}) = \big(1 - F(\varphi^{-1}(s_i))\big)^{I-1} \cdot U(s_i, \theta_i \mid \mathcal{A}).$$

Proposition B.1. *Each well-behaved auction $\mathcal{A}$ has a unique symmetric equilibrium. This equilibrium is increasing, differentiable in the bidder's type θ, and is characterized as the unique solution to the following differential equation:*

$$\varphi'(\theta) = \frac{(I-1)f(\theta)}{1 - F(\theta)} \cdot \frac{U(\varphi(\theta), \theta \mid \mathcal{A})}{U_1(\varphi(\theta), \theta \mid \mathcal{A})}, \tag{B.2}$$

subject to the boundary condition $U(\varphi(\overline{\theta}), \overline{\theta} \mid \mathcal{A}) = 0$.

Proof. We begin by establishing monotonicity of the equilibrium bidding strategy using Topkis's Monotonicity Theorem (Topkis, 1978). Consider a seller of type θ facing an arbitrary strategy profile $\sigma : [\underline{\theta}, \overline{\theta}] \to \mathbb{R}$ of the other $I - 1$ opponents, and let $P(s)$ be the probability that a bid s wins against the opponents' bids. Since the opponents' types are independent of θ, $P(s)$ does not depend on type.

The expected utility of bidding s is

$$V(s, \theta) = P(s) \cdot U(s, \theta \mid \mathcal{A}).$$

By assumption, $U(s,\theta \mid \mathcal{A})$ is increasing in s, decreasing in θ, differentiable, and log-supermodular on Π_+. Since $\log V(s,\theta) = \log P(s) + \log U(s,\theta \mid \mathcal{A})$ and $\log P(s)$ depends only on s, the restriction of $\log V$ to Π_+ preserves supermodularity in (s,θ). By Topkis' Monotonicity Theorem, any selection $\varphi(\theta) = \arg\max_s V(s,\theta)$ is weakly increasing in θ.

If $\varphi(\theta)$ were constant over an interval, then the highest type in that interval could profitably increase her bid, raising her probability of winning and expected utility. Therefore, $\varphi(\theta)$ must be strictly increasing.

Having established that best responses are strictly increasing, we now characterize the symmetric equilibrium by deriving the differential equation satisfied by $\varphi(\cdot)$. Assume all sellers use the same strictly increasing, differentiable strategy $\varphi(\cdot)$. For type θ, the probability of winning with a bid s is $(1 - F(\varphi^{-1}(s)))^{I-1}$, so the expected utility is

$$V(s,\theta) = (1 - F(\varphi^{-1}(s)))^{I-1} \cdot U(s,\theta \mid \mathcal{A}).$$

In equilibrium, the optimal bid is $s = \varphi(\theta)$, and the first-order condition requires

$$\left.\frac{\partial V}{\partial s}\right|_{s=\varphi(\theta)} = 0.$$

Differentiating,

$$\begin{aligned}
\frac{\partial V}{\partial s} &= \frac{\partial}{\partial s}\Big[(1 - F(\varphi^{-1}(s)))^{I-1}\Big] U(s,\theta \mid \mathcal{A}) \\
&\quad + (1 - F(\varphi^{-1}(s)))^{I-1} U_1(s,\theta \mid \mathcal{A}) \\
&= -(I-1)(1 - F(\varphi^{-1}(s)))^{I-2} f(\varphi^{-1}(s)) \frac{1}{\varphi'(\varphi^{-1}(s))} U(s,\theta \mid \mathcal{A}) \\
&\quad + (1 - F(\varphi^{-1}(s)))^{I-1} U_1(s,\theta \mid \mathcal{A}).
\end{aligned}$$

Evaluating at $s = \varphi(\theta)$ yields

$$(I-1)(1-F(\theta))^{I-2} f(\theta) \frac{1}{\varphi'(\theta)} U(\varphi(\theta), \theta \mid \mathcal{A}) = (1-F(\theta))^{I-1} U_1(\varphi(\theta), \theta \mid \mathcal{A}).$$

Rearranging gives

$$\varphi'(\theta) = \frac{(I-1)f(\theta)}{1-F(\theta)} \cdot \frac{U(\varphi(\theta), \theta \mid \mathcal{A})}{U_1(\varphi(\theta), \theta \mid \mathcal{A})}.$$

The boundary condition $U(\varphi(\overline{\theta}), \overline{\theta} \mid \mathcal{A}) = 0$ follows from the requirement that the highest type is indifferent to participating. Since the probability of winning with a bid corresponding to the highest type is zero, $P(\varphi(\overline{\theta})) = 0$, and all types must earn non-negative expected utility, it follows that $U(\varphi(\overline{\theta}), \overline{\theta} \mid \mathcal{A}) \leqslant 0$. If the inequality were strict, the highest type could profitably raise her bid slightly, increasing her winning probability and expected utility. Thus, the boundary condition must bind.

Finally, global optimality follows because maximizing $V(s,\theta)$ is equivalent to maximizing $\log V(s,\theta)$, and $\log U(s,\theta \mid \mathcal{A})$ has strictly increasing differences on Π_+. Hence, the first-order condition is sufficient, and φ is the unique symmetric equilibrium bidding strategy. □

B.4 Utility equivalence

This section characterizes utility equivalence across different scoring auctions. We begin by introducing the notion of homogeneous type effects, which provides the key condition for such equivalence.

Definition B.2. *Two scoring auctions $\mathcal{A}$ and $\mathcal{A}_*$ with respective equilibrium bidding strategies φ and φ_* satisfy **homogeneous type effects** if they have identical marginal effects of the seller's type on expected utility at their respective equilibrium scores:*

$$U_2(\varphi(\theta), \theta \mid \mathcal{A}) = U_2(\varphi_*(\theta), \theta \mid \mathcal{A}_*), \quad \text{for each } \theta \in \Theta,$$

where $U_2(s,\theta \mid \mathcal{A})$ denotes the partial derivative of $U(s,\theta \mid \mathcal{A})$ with respect to θ.

This condition ensures that the sensitivity of expected utility to changes in a seller's private information is identical across both auction formats. The following result shows that, when this condition holds, the auctions yield the same expected utility for each seller type.

Proposition B.2. *Fix two scoring auctions $\mathcal{A}$ and $\mathcal{A}_*$ with respective equilibrium bidding strategies φ and φ_*. The auctions induce the same expected utility for every seller type if and only if they have homogeneous type effects.*

Proof. Let $\mathcal{U}(\theta) = U(\varphi(\theta), \theta \mid \mathcal{A})$ and $\mathcal{U}_*(\theta) = U(\varphi_*(\theta), \theta \mid \mathcal{A}_*)$ denote the equilibrium expected utilities of type θ in auctions $\mathcal{A}$ and $\mathcal{A}_*$, respectively.

We first show the "if" direction. Suppose $\mathcal{U}(\theta) = \mathcal{U}_*(\theta)$ for all θ. Then $\mathcal{U}'(\theta) = \mathcal{U}_*'(\theta)$ for all θ. Differentiating $\mathcal{U}(\theta) = U(\varphi(\theta), \theta \mid \mathcal{A})$ with respect to θ gives

$$\begin{aligned}
\mathcal{U}'(\theta) &= U_1(\varphi(\theta), \theta \mid \mathcal{A}) \cdot \varphi'(\theta) + U_2(\varphi(\theta), \theta \mid \mathcal{A}) \\
&= -\frac{(I-1)f(\theta)}{1-F(\theta)} U(\varphi(\theta), \theta \mid \mathcal{A}) + U_2(\varphi(\theta), \theta \mid \mathcal{A}) \\
&= -\frac{(I-1)f(\theta)}{1-F(\theta)} \mathcal{U}(\theta) + U_2(\varphi(\theta), \theta \mid \mathcal{A}),
\end{aligned}$$

where the second equality follows from Proposition B.1. An analogous expression holds for $\mathcal{U}_*'(\theta)$. Since $\mathcal{U}(\theta) = \mathcal{U}_*(\theta)$ and $\mathcal{U}'(\theta) = \mathcal{U}_*'(\theta)$, it follows that

$$U_2(\varphi(\theta), \theta \mid \mathcal{A}) = U_2(\varphi_*(\theta), \theta \mid \mathcal{A}_*), \tag{B.3}$$

which establishes that the auctions satisfy homogeneous type effects.

For the "only if" direction, assume the auctions satisfy homogeneous type effects, so Equation (B.3) holds. Since both auctions share the boundary condition $U(\varphi(\overline{\theta}), \overline{\theta} \mid \mathcal{A}) = U(\varphi_*(\overline{\theta}), \overline{\theta} \mid \mathcal{A}_*) = 0$, integrating the equal derivatives from $\overline{\theta}$ to θ implies that $\mathcal{U}(\theta) = \mathcal{U}_*(\theta)$ for all θ. $\square$

Definition B.2 requires that the marginal effect of the seller's type θ be identical across both auction formats. In essence, the marginal sensitivity of the winning payoff—in certainty-equivalent terms—to private information must coincide under $\mathcal{A}$ and $\mathcal{A}_*$. Importantly, this equivalence is required only at the equilibrium bids $\varphi(\theta)$ and $\varphi_*(\theta)$.

While the condition of homogeneous type effects may be difficult to verify in some settings, the next subsection provides two sufficient conditions for utility equivalence that are straightforward to check.

B.4.1 Separability

We now focus on a simple sufficient condition on the primitives that ensures utility equivalence. Fix a well-behaved scoring auction $\mathcal{A}$. Say that $\mathcal{A}$ is **separable** if there exists a continuous, strictly increasing mapping $R : \mathbb{R} \to \mathbb{R}$ such that

$$U(s, \theta \mid \mathcal{A}) = u\big(R(s) - k(\theta)\big), \quad \text{for each } \theta \in \Theta \text{ and each } s \in S.$$

An auction is separable when all types agree that the certainty-equivalent value of each score s is $R(s)$. This property implies that $\mathcal{A}$ is, in a sense, equivalent to a cash auction in which bidders face only fixed costs $k(\theta)$ and bid over deterministic payments $R(s)$ that are independent of type.

Proposition B.3. *Fix two well-behaved auctions $\mathcal{A}$ and $\mathcal{A}_*$. If both are separable, then the equilibrium expected utility of each type θ is the same in $\mathcal{A}$ and in $\mathcal{A}_*$.*

For the proof, see Corollary 1 in the main text.

Assumption 1. *The environment satisfies the following:*

The utility function $u(\cdot)$ has null wealth effects.[39]

The input distribution $\mu(\theta)$ and marginal costs $\mathbf{c}(\theta)$ are constant in θ.

[39]For instance, a CARA utility function with parameter γ, $u(\pi) = 1 - \exp(-\gamma\pi)$, satisfies this property.

Theorem B.1. *If the environment satisfies Assumption 1, then all rich scoring auctions are well-behaved and separable. Consequently, all rich scoring auctions yield the same utility to each type.*

Proof. By Proposition B.3, it suffices to show that each rich scoring auction is well-behaved and separable. Fix an alternative $\boldsymbol{b} \in B(s)$ for some score s and let $\mathrm{CE}(b)$ denote the certainty-equivalent payment of the lottery $\Pi(\boldsymbol{q}, s, b) - \boldsymbol{q} \cdot \boldsymbol{c}$, which satisfies

$$u\big(\mathrm{CE}(b)\big) = \mathop{\mathbb{E}}_{\boldsymbol{q}\sim\mu}\Big[u\big(\Pi(\boldsymbol{q}, s, b) - \boldsymbol{q} \cdot \boldsymbol{c}\big)\Big].$$

Since u is continuous and strictly monotone, the certainty equivalent $\mathrm{CE}(b)$ is well defined and unique for each alternative $b \in B(s)$. Define $R(s) = \max_{b\in B(s)} \mathrm{CE}(b)$. Because there are null wealth effects, each type θ solves

$$U(s, \theta) = \max_{b\in B(s)} u\big(\mathrm{CE}(b) - k(\theta)\big) = u\big(R(s) - k(\theta)\big).$$

This establishes that the scoring auction is well-behaved and satisfies separability. □

Theorem B.1 yields two significant simplifications of the analysis. First, sellers are completely indifferent across auction formats. This indifference applies not only to scaling and cash auctions discussed in the main text, but also to security-bid auctions, as in a procurement adaptation of DeMarzo, Kremer, and Skrzypacz (2005b). As a result, an analyst seeking a Pareto improvement should select the auction that minimizes the buyer's expected payment. Second, because sellers are risk averse, the buyer minimizes expected payments by choosing the auction format that offers greater insurance.

B.4.2 Competitive Environments

This subsection shows that utility equivalence holds in competitive environments as the number of sellers is large. We consider a continuum (unit mass) of sellers whose types θ are independently and identically distributed according to an atomless distribution F. In addition, we drop the requirement that all types share the same utility function, allowing for the case in which each type θ has a Bernoulli utility function u_θ. Throughout this section, we restrict attention to rich and well-behaved auctions, ensuring that $U(s, \theta \mid \mathcal{A})$ is strictly increasing in s.

Theorem B.2. *Suppose there is a continuum (unit mass) of sellers with i.i.d. types drawn from an atomless distribution. Let $\mathcal{A}$ be a rich and well-behaved scoring auction. Then there exists a Bayesian equilibrium in which each type receives zero expected utility. Moreover, in every symmetric Bayesian equilibrium of $\mathcal{A}$, each type $\theta \in \Theta$ obtains zero expected utility.*

Proof. We begin with existence. Define

$$S_0(\theta) := \left\{ s \in [\underline{s}, \overline{s}] : U(s, \theta) = 0 \right\}$$

as the set of scores that yield zero utility for type θ. Consider a strategy profile φ such that $\varphi(\theta) \in S_0(\theta)$ for all $\theta \in \Theta$. Note that φ is continuous and increasing. The probability of winning with a score s is 1 for $s < \varphi(0)$ and 0 for $s \geqslant \varphi(0)$.

Fix a type $\theta \in \Theta$. Any deviation that results in a positive winning probability requires submitting a score $s < \varphi(0)$, which yields negative utility. Hence no type has an incentive to deviate, and this profile is a Bayesian equilibrium.

We now show that in every equilibrium, each type receives zero expected utility. Let $\sigma : \Theta \to [\underline{s}, \overline{s}]$ be a symmetric Bayesian equilibrium and define

$$s_* := \inf \left\{ s : s \in \bigcup_{\theta \in \Theta} \operatorname{supp}(\sigma(\theta)) \right\},$$

the infimum of the scores selected under σ. With a continuum of sellers, the winning score is s_*. Thus, any score $s < s_*$ wins with probability one, while any $s > s_*$ wins with probability zero.

Fix a type $\theta \in \Theta$ and let $\mathcal{U}(\theta)$ be its equilibrium expected utility. Suppose $\mathcal{U}(\theta) >$

0. Then the equilibrium score of type θ must satisfy $s(\theta) < s_*$; otherwise, the winning probability would be zero. But for any $s' \in (s(\theta), s_*)$, the winning probability remains one, and since $U(\cdot, \theta)$ is strictly increasing on U_+, we obtain $U(s', \theta) > U(s(\theta), \theta)$, contradicting optimality. Hence $\mathcal{U}(\theta) \leqslant 0$ for all θ.

Conversely, any type can secure zero expected utility by bidding any $\tilde{s} \geqslant s_*$, which yields a winning probability of zero. Thus $\mathcal{U}(\theta) \geqslant 0$. Combining the two inequalities implies $\mathcal{U}(\theta) = 0$ for each $\theta \in \Theta$. $\square$

B.5 Second-price variants

Following Che (1993), this section analyzes second-price variants of scoring auctions in environments that satisfy separability (i.e., Assumption 1). In the second-price variant, while the seller with the lowest score wins the auction, the winner is not constrained by their original alternative nor the second-lowest bidder's alternative. Rather, the winning seller has the flexibility to choose any alternative that generates a score equal to the second-lowest score. This feature ensures that the winner can improve the outcome while satisfying the second-lowest score requirement.

Building on our theoretical framework, we establish utility equivalence between second-price scoring auctions and their corresponding auxiliary auction counterparts. In the second-price auxiliary auction, sellers submit certainty equivalents of the induced lottery, with the seller submitting the lowest score winning but receiving payment according to the second-lowest score.

Theorem B.3 (Utility Equivalence for Second-Price Scoring Auctions)**.** *Under separability (Assumption 1), the second-price variants of all well-behaved scoring auctions yield identical expected utility for every seller type. That is, for any two well-behaved scoring auctions $\mathcal{A}$ and $\mathcal{A}_*$ with respective second-price equilibrium strategies φ and φ_*, we have*

$$U(\varphi(\theta), \theta \mid \mathcal{A}) = U(\varphi_*(\theta), \theta \mid \mathcal{A}_*), \quad \textit{for each } \theta \in \Theta.$$

Proof. Under assumption 1, there is some strictly increasing mapping $R : \mathbb{R} \to \mathbb{R}$

such that $U(s, \theta \mid \mathcal{A}) = u(R(s) - k(\theta))$. (See Theorem B.1.) Likewise, there is some strictly increasing mapping $R^* : \mathbb{R} \to \mathbb{R}$ such that $U(s^*, \theta \mid \mathcal{A}_*) = u(R^*(s^*) - k(\theta))$. In the second-price auxiliary auction, sellers submit certainty equivalents of the induced lottery, with the seller submitting the lowest score winning but receiving the second-lowest certainty equivalent. Hence, it is dominant for each type θ to select a score s in auction $\mathcal{A}$ that delivers a certainty equivalent payment equal to $k(\theta)$, i.e., a score such that $R(s) = k(\theta)$. Likewise, it is dominant for each type θ to select a score s^* in auction $\mathcal{A}_*$ that delivers a certainty equivalent payment equal to $k(\theta)$, i.e., a score such that $R^*(s^*) = k(\theta)$. As a result, each type θ selects a score s in $\mathcal{A}$ and a score s^* in $\mathcal{A}_*$ such that $R(s) = R^*(s^*) = k(\theta)$, yielding $U(\varphi(\theta), \theta \mid \mathcal{A}) = U(\varphi_*(\theta), \theta \mid \mathcal{A}_*)$. $\square$

Remark B.1. *Under separability, utility equivalence holds within first-price formats and within second-price formats. However, no such utility equivalence exists between first-price and second-price variants of these auctions.*

Payment minimization Drawing on classical results from Maskin and Riley (1984); Krishna (2009), sellers submit lower equilibrium scores in first-price auxiliary auctions compared to second-price auxiliary auctions. The first-price format effectively induces sellers to purchase insurance against the possibility of losing the auction. By leveraging our equivalence results, this pattern extends to scoring auctions: sellers in first-price formats request lower certainty equivalents than in second-price formats, resulting in lower scores and reduced expected payments from the buyer. A benevolent seller prioritizing buyer welfare may prefer implementing second-price variants.

B.6 Beyond Utility Equivalence

The results in the main text depend on the conditions that guarantee utility equivalence. When these conditions fail, the novel Hicksian decomposition methodology introduced in the paper is no longer applicable. Three important instances in which utility equivalence breaks down are discussed below.

Wealth effects Utility equivalence does not hold for utility functions that exhibit wealth effects, such as CRRA and HARA, where risk aversion depends on the seller's wealth level. In these cases, the certainty equivalent payment depends on the agent's wealth, so the separability property fails and the expected utilities delivered by different auction formats change depending on the scoring rule. Nevertheless, simulations show that these differences in utilities are typically small, with outcomes remaining close to those under separability. Consequently, using expected quantities as weights approximately minimizes the government's expected payment. Simulation results showing the approximate optimality of using expected quantities as weights are available upon request.

Heterogeneous marginal costs Utility equivalence can break down abruptly in the presence of even small heterogeneous marginal costs. In contrast to wealth effects, which cause only gradual deviations, cost heterogeneity can radically change which seller types possess a comparative advantage. When sellers differ in their marginal costs across tasks, adjusting the scoring weights may entirely change the identity of the winning bidder. A type that dominates under one set of weights may become uncompetitive under another, as the scoring rule alters the relative appeal of different bidding options. Consequently, even slight differences in marginal costs can result in substantial changes in sellers' expected utilities.

Heterogeneous distributions Utility equivalence breaks down when seller types face different demand distributions, fundamentally altering the auction dynamics. When higher-cost (inefficient) types face distributions that generate higher expected quantities, these types develop stronger incentives to win the project. This creates a perverse effect: inefficient types bid more aggressively than efficient ones, potentially capturing the project despite their higher costs. Unlike cash auctions, where efficient types typically prevail, scaling auctions can induce the most inefficient types to win. This reversal occurs because the scoring mechanism directly rewards expected quantity, aligning the auction's incentives with the inefficient types' comparative advantage.

These cases highlight important avenues for further research. Our framework and Hicksian decomposition apply systematically in two classes of environments: competitive settings or separable environments—in which utility equivalence holds. Extending this methodology to settings with wealth effects or heterogeneous costs remains an open challenge, and our analysis offers a building block for the analysis of these more general environments.