Risk Robust Mechanism Design for a Prospect Theoretic Buyer

Siqi Liu[*] J. Benjamin Miller[†] Alexandros Psomas[‡]

Abstract

Consider the revenue maximization problem of a risk-neutral seller with m heterogeneous items for sale to a single additive buyer, whose values for the items are drawn from known distributions. If the buyer is also risk-neutral, it is known that a simple and natural mechanism, namely the better of selling separately or pricing only the grand bundle, gives a constant-factor approximation to the optimal revenue. In this paper we study revenue maximization without risk-neutral buyers. Specifically, we adopt cumulative prospect theory, a well established generalization of expected utility theory.

Our starting observation is that such preferences give rise to a very rich space of mechanisms, allowing the seller to extract arbitrary revenue. Specifically, a seller can construct extreme lotteries that look attractive to a mildly optimistic buyer, but have arbitrarily negative true expectation. Therefore, giving the seller absolute freedom over the design space results in absurd conclusions; competing with the optimal mechanism is hopeless. Instead, in this paper we study four broad classes of mechanisms, each characterized by a distinct use of randomness. Our goal is twofold: to explore the power of randomness when the buyer is not risk-neutral, and to design simple and attitude-agnostic mechanisms—mechanisms that do not depend on details of the buyer's risk attitude—which are good approximations of the optimal in-class mechanism, tailored to a specific risk attitude. Our main result is that the same simple and risk-agnostic mechanism (the better of selling separately or pricing only the grand bundle) is a good approximation to the optimal non-agnostic mechanism within three of the mechanism classes we study.

1 Introduction

Expected utility theory (EUT) has long reigned as the prevailing model of decision making under uncertainty. However, a substantial body of evidence, including the famous Allais paradox [1], shows that most people make choices that violate this theory. Cumulative prospect theory (CPT) (Tversky and Kahneman [33]) is arguably the most prominent alternative. A key element of this theory is a non-linear transformation of cumulative probabilities by a *probability weighting function*. This transformation can model a person's tendency towards optimism or pessimism.[1] On the other hand, as mechanism designers we use randomization as an important tool in optimizing our objective, typically (and crucially) assuming that agents make choices according to the tenets of expected utility theory. While we have vastly deepened our understanding of mechanism design under this assumption, it is essential to study empirically validated models of human decision-making.

In this paper we study the revenue-maximization problem of a risk-neutral seller with m heterogeneous items for sale to a single, additive buyer with cumulative prospect theory preferences.

[*]UC Berkeley
[†]UWMadison
[‡]Carnegie Mellon University

[1]As we discuss below, real-world attitudes are not merely "optimistic" or "pessimistic", but such simplistic attitudes are easily and naturally captured by this model.

Our goal is to design simple mechanisms that are agnostic to the underlying probability weighting function of the buyer, yet achieve a good approximation to the revenue of the optimal mechanism *tailored* to this weighting function. To understand our results in context, we begin by briefly reviewing cumulative prospect theory.

1.1 Prospect Theory Basics

In full generality, cumulative prospect theory (CPT) asserts that preferences are parameterized by a reference point (or status quo) r, a value function U that maps (deterministic, i.e. certain) outcomes into *utils* (or dollars), and two probability weighting functions, w^+ and w^-, for weighting the cumulative probabilities of positive and negative outcomes (relative to r). Throughout the paper, and like most works in mechanism design, we assume linear utility for money: $U(x) = x$. That is, our agents have value 1 for \$1 and value 1000 for \$1000. What remains, then, are the weighting functions w^+ and w^- and the reference point r. By taking $r = 0$ and the weighting functions to be the identity function, one recovers expected utility theory; thus, CPT generalizes expected utility theory.

To gain an understanding of this theory, consider first a simple event E which occurs with probability $\frac{1}{2}$, and assume that $r = 0$. Suppose that E corresponds to an agent receiving value 10; if E does not occur, the agent receives nothing. A risk-neutral agent would value this potential income at $10 \cdot \Pr[E] = 5$. An optimistic agent, overestimating the possibility of receiving 10, might value E at slightly more than 5, whereas a pessimistic agent might value it at slightly less. CPT uses a weighting function w^+ which modifies probabilities of positive (with respect to r) outcomes: the agent values event E at $10 \cdot w^+(\Pr[E])$. Then $w^+(x) > x$ corresponds to optimism, and $w^+(x) < x$ corresponds to pessimism. CPT captures much more complex behavior than merely optimism and pessimism. For example, in experiments (e.g. [33, 8]), subjects tend to overweight extreme events: in a sense, people are optimistic about very good outcomes and pessimistic about very bad outcomes. This sort of behavior can be readily captured by CPT, and as it turns out, it suggests inverse-S-shaped weighting functions.

In general, the event of interest might correspond to a positive or negative outcome. For example, E might correspond to the agent *losing* value 10. In that case, we expect the optimistic agent to *underweight* the probability of E occurring. For this reason, CPT models probability weighting for gains and losses with functions w^+ and w^-, respectively. Furthermore, when the random variable is supported on multiple non-zero values, applying w^+ (or w^-) directly to the probability of each event leads to violations of first-order stochastic dominance. For this reason, Quiggin [28] proposed to weight the cumulative distribution function, rather than the probability mass function; hence *cumulative* prospect theory.

Our interest here is highlighting the effects of nonlinear probability weighting. We will therefore focus on a special case of cumulative prospect theory, namely *rank dependent utility theory* (RDUT). This theory is rich enough to explain a number of known violations of expected utility theory, e.g., the Allais paradox [29], general enough to include expected utility theory as a special case, while at the same time simple enough to be mathematically tractable. This theory is equivalent to the following assumption.

Assumption 1. *For all* $p \in [0, 1]$, $w^-(p) = 1 - w^+(1 - p)$ *(Quiggin [28]).*

Assumption 1 allows us to rank all the outcomes from worst to best, independent of whether they are gains or losses, and weight their probabilities with a single weighting function $w(x)$. Furthermore, it makes the reference point r irrelevant. Chawla et al. [11] have previously studied the same model, giving a class of mechanisms which optimally sell a single item to a pessimistic buyer. However, they restrict themselves to convex weighting functions. Here we study general weighting functions and multi-item auctions. We postpone more details about rank dependent utility theory until Section 2, and refer the reader to Appendix A for what the expected utility of a general CPT agent (that is, without Assumption 1) for even a simple lottery looks like.

1.2 Our Results

Our starting point is the observation that even very mild probability weighting gives rise to rich seller behavior, which allows the seller to extract unbounded revenue. Specifically, we show that under assumptions satisfied by most weighting functions in the literature, the seller can design a bet that has arbitrarily negative (risk-neutral) expectation, but looks attractive to a RDUT buyer. This bet can be easily turned into an auction for selling any number of items by giving the items for free if and only if the buyer takes the bet. Similar behavior has been observed before this work for more general models, e.g. by Azevedo and Gottlieb [2] and Ebert and Strack [15].

In light of these negative results for arbitrary buyer-seller interaction, we focus our attention to specific classes of mechanisms, imposing various restrictions on the mechanism's description and implementation. These restrictions are not onerous: when offered to a risk-neutral buyer, two of the classes are equivalent to the class of all mechanisms, and another is equivalent to all deterministic mechanisms. Our restrictions thus serve to isolate particular uses of randomization and to illustrate the various effects RDUT preferences have on mechanism design.

The first class we consider is that of *deterministic price* mechanisms, which we denote $\mathcal{C}_{dp}$. Here, the seller offers a menu of (possibly correlated) distributions over the items, each at a fixed price. The buyer may pay the price for a distribution, after which she receives a draw from the distribution. To bypass some technical barriers, we also consider a special case of this class, *nested deterministic price* mechanisms, or $\mathcal{C}_{ndp}$, which impose certain constraints on the distributions over items in a menu. These constraints are very mild (for example they are always satisfied by independent distributions) and are without loss of generality for a risk-neutral buyer. Next, we consider the class of *deterministic allocation* mechanisms, $\mathcal{C}_{da}$, where the mechanism deterministically allocates a bundle of items for a possibly randomized, non-negative payment. $\mathcal{C}_{da}$ is equivalent to deterministic mechanisms for a risk neutral buyer. Finally, we consider a multi-item generalization of the single-item class of mechanisms that is optimal for convex weighting functions (as shown by Chawla et al. [11]). We call this class *binary-lottery* mechanisms and denote it by $\mathcal{C}_b$.

Our main result is that, for classes $\mathcal{C}_{ndp}$, $\mathcal{C}_{da}$ and $\mathcal{C}_b$, a single simple mechanism, agnostic to the underlying weighting function, gives a good approximation on the revenue of the *optimal* in-class mechanism *tailored* to w. That mechanism is the better of selling every item separately at a fixed price (henceforth SRev) and selling the grand bundle as a single item at a fixed price (henceforth BRev), which is a valid mechanism in all classes considered. Furthermore, this mechanism is deterministic, which implies that its expected revenue is the same for all weighting functions w, and only depends on the buyer's value distribution $\mathcal{D}$. Our proof is by relating the revenue of each class of mechanisms to the revenue obtainable from a risk-neutral buyer via any mechanism, combined with a result of Babaioff et al. [3], which shows that $\max\{\text{SRev}, \text{BRev}\}$ is a constant approximation to this risk-neutral revenue. For $\mathcal{C}_{dp}$ our understanding is partial; we show that $\max\{\text{SRev}, \text{BRev}\}$ approximates the optimal, risk non-agnostic $\mathcal{C}_{dp}$ auction within a doubly exponential in the number of items factor. This, of course, implies a constant approximation for a constant number of items (in fact, for two items we can show an approximation factor of 2 for just SRev), but we leave it as an open problem whether a constant approximation is possible for the general case. All our results can be extended to a unit-demand and additive up to a downward closed constraint buyer by paying an extra factor of 4 and 31.1, respectively, using the results of Chawla and Miller [10].

Intuitively, the difficulty with analyzing mechanisms for RDUT buyers (and especially optimal mechanisms) is that, given a mechanism, we cannot generally argue about how much a buyer type t values the menu item purchased by a type t'. This is especially the case for general deterministic price mechanisms, where allocations over items could be arbitrarily correlated. This, in turn, prevents us from using basic "simulation arguments": starting from an auction $\mathcal{M}$, manipulate the allocation rule and pricing rule to get a different auction $\mathcal{M}'$. Such arguments are very useful in getting meaningful upper bounds on the optimal revenue. For example, Hart and Nisan [22] upper bound the optimal revenue from a product distribution, $\text{Rev}(\mathcal{D} \times \mathcal{D}')$, by $\text{Rev}(\mathcal{D}) + \text{VAL}(\mathcal{D}')$[2] using such an argument, where they give a concrete auction for $\mathcal{D}$ by manipulating the allocation and payment rule of the optimal auction for $\mathcal{D} \times \mathcal{D}'$. Similar "marginal mechanism" arguments are crucial in many works that give simple and approximately optimal mechanisms for additive buyers, e.g. Babaioff et al. [3], Yao [37], Li and Yao [24]; for example, the so-called core-tail decomposition technique depends on such arguments. On the other hand, the recently developed Lagrangian duality based approach (Cai et al. [7], Fu et al. [20], Cai and Zhao [6], Devanur and Weinberg [12], Eden et al. [16, 17], Liu and Psomas [25]) also seems to fail here. This technique has been successful in getting benchmarks in a number of settings, by giving a solution to the dual of the mathematical program that computes the optimal auction. To the best of our knowledge, all works that use this technique start from a linear program. Here, the mathematical program for the optimal, risk non-agnostic auction is not even convex. Even though in theory only weak duality is necessary for this technique to work, we haven't been successful in applying it to our problem.

Before we proceed, we mention here another reasonable approach to model robustness with respect to risk. Find the mechanism $\mathcal{M}$ that maximizes (over all mechanisms) the seller's revenue in the worst case with respect to the weighting function w (similar to recent results of Carroll [9], Gravin and Lu [21] for robustness with respect to correlation). In this scenario, we observe (see Appendix B) that the optimal mechanism is the optimal deterministic mechanism, so we get that $\max\{\text{SRev}, \text{BRev}\}$ is a good approximation to the optimal revenue by a trivial reduction to the risk neutral-buyer setting. To see why this is the case, notice that $w(x)$ could take the value one for all x except $x = 0$. In this case the buyer is extremely optimistic; her utility for a random variable X is equal to her utility for her favorite outcome. We argue that randomizing only hurts the seller, since doing so decreases both the expected utility of the buyer, and the expected revenue of the seller.

1.3 Related Work and Roadmap

Prospect theory was originally defined by Kahneman [23] but, though successful in explaining experimentally observed behavior, it suffered from a number of weaknesses, namely violations of first-order stochastic dominance between random variables. Several works (Weymark [35], Quiggin [28], Yaari [36], Schmeidler [32]) proposed solutions to these issues, resulting in cumulative prospect theory (Tversky and Kahneman [33]). Next to expected utility theory, cumulative prospect theory is likely the best studied theory of decision-making under uncertainty. We refer the reader to the book of Wakker [34] for a thorough exposition of the model. Also see Machina [26] for a survey of non-EUT models.

Although widely studied in behavioral economics, prospect theory has received much less attention in the game theory and mechanism design literature. Our work is most closely related to that of Chawla et al. [11], who study optimal and robust mechanisms for a single buyer and a single item. Their work, unlike ours, places much stronger assumptions on the weighting function:

[2] $\mathcal{D}$ and $\mathcal{D}'$ here are distribution over m_1 and m_2 items, respectively. $\text{VAL}(\mathcal{D}') = \sum_{j \in [m_2]} \mathbb{E}[\mathcal{D}'_j]$, i.e. the total expected sum of values from items in $\mathcal{D}'$.

namely, they assume convexity (which in turn implies $w(x) \leq x$). In this paper we consider general weighting functions, but restrict the mechanism design space. Further afield, Easley and Ghosh [14] study contract design in a crowdsourcing setting with a prospect-theoretic model of workers. Fiat and Papadimitriou [18] demonstrate that equilibria may not exist in two-player games when players have prospect-theoretic preferences. Dughmi and Peres [13] and Fu et al. [19] study mechanism design with risk-averse agents in a setting where risk-averse behavior is represented by a concave utility function, while more recently, in a similar setting, Nikolova et al. [27] study optimal mechanisms for risk-loving agents.

Our main result is that the better of selling separately and selling the grand bundle is a risk robust approximation to the optimal revenue. The approximation ratio of this mechanism has been studied extensively for risk-neutral buyers having a large class of valuations [3, 5, 31, 7, 10, 5]. Our result relies on this work, but our techniques are very different.

Roadmap. Section 2 poses our model and some preliminaries. We discuss the limits of our model in Section 3, and show that if the seller is allowed to use an arbitrary mechanism, then he can extract arbitrarily large revenue. In Section 4 we formally define the mechanism classes considered in this paper, which we proceed to analyze in Sections 5 (deterministic price mechanisms), 6 (deterministic allocation mechanisms) and 7 (binary lottery mechanisms).

2 Preliminaries

A risk-neutral seller, whose aim is to maximize revenue, is auctioning off m items to a single buyer with cumulative prospect theory preferences. The value of the buyer for item i is v_i, and is distributed according to a known distribution $\mathcal{D}_i$. We assume that the item distributions are independent, and denote the joint distribution by $\mathcal{D}$. We first go over the buyer's preference model in detail, and then formulate our mechanism design problem.

Weighted Expectation. In this paper we focus on a special case of cumulative prospect theory, *rank dependent utility theory*. In rank dependent utility theory a weighting function w distorts cumulative probabilities (Quiggin [28]). The weighting function w satisfies the following properties: (1) $w : [0, 1] \to [0, 1]$, (2) w is non-decreasing, (3) $w(0) = 0$ and $w(1) = 1$. We use the notation $\mathcal{I}$ to indicate the risk-neutral weighting function; that is $\mathcal{I}(x) = x$. For a random variable Z over k outcomes, where the i-th outcome occurs with probability p_i and gives utility u_i, and $u_i \leq u_{i+1}$, an agent with weighting function w has expected utility

$$\mathbb{E}_w\left[Z\right] = \sum_{i=1}^{k-1} u_i \left(w\left(\sum_{j=i}^{k} p_j\right) - w\left(\sum_{j=i+1}^{k} p_j\right)\right) + u_k w(p_k) = u_1 + \sum_{i=2}^{k} (u_i - u_{i-1}) \cdot w\left(\sum_{j=i}^{k} p_j\right).$$

The intuitive interpretation (for the latter expression) is that the agent always gets utility u_1. Then, the event that the agent gets an additional utility of at least $u_2 - u_1$ occurs with probability $1 - p_1 = \sum_{j=2}^{k} p_j$ (which is weighted by the function w). The agent gets an additional utility of at least $u_3 - u_2$ with probability $\sum_{j=3}^{k} p_j$, and so on. We note that this definition makes no assumption about the sign of u_i; that is, the u_is can be positive (corresponding to gains) or negative (corresponding to losses). We will also use the following, equivalent definition.

Definition 1. ***(Weighted expectation, probability version).*** *Let Z be a random variable supported in $(-\infty, \infty)$ with cumulative distribution function F_Z. Then the* weighted expectation *of Z with respect to the weighting function w is*

$$\mathbb{E}_w[Z] \;=\; -\int_{-\infty}^{0} \left[1 - w\big(1 - F_Z(z)\big)\right] dz \;+\; \int_{0}^{\infty} w\big(1 - F_Z(z)\big) dz. \tag{1}$$

Before continuing our exposition, we illustrate the flexibility and power of the model through some examples. The weighting function can be used also to pick out various statistics of interest, as the next two examples illustrate.

Example 1. *Suppose w is given by*

$$w(x) = \begin{cases} 0 & x \in [0, \frac{1}{2}) \\ 1 & x \in [\frac{1}{2}, 1]. \end{cases}$$

Then the w-weighted expectation of any random variable is equal to its median.

In general, of course, Example 1 could be modified to pick out any quantile; simply change $\frac{1}{2}$ in the definition of w to some $\tau \in [0, 1]$. As special cases, the *maxmin* choice rule, in which the agent evaluates a distribution according to the worst-case outcome, corresponds to picking $\tau = 1$, and the *maxmax* choice rule corresponds to $\tau = 0$.[3]

Example 2. *Suppose $w(x)$ is given by*

$$w(x) = \begin{cases} \theta & x \in [0, 1) \\ 1 & x = 1, \end{cases}$$

for some $\theta \in (0, 1)$. Then the w-weighted expectation is equal to the weighted average of the highest and lowest outcomes in the support. E.g., if Z is supported on $[L, H]$, then $\mathbb{E}_w[Z] = (1-\theta)L + \theta H$.

We note that Examples 1 and 2 are among the utility functions studied by Fiat and Papadimitriou [18]; they show that games need not possess any Nash equilibrium when the players seek to maximize the corresponding weighted expectations.

Properties of weighted expectation. Weighted expectation in this model satisfies the following properties; proofs can be found in Appendix C.

Lemma 1. *For any weighting function w, any random variable Z and any $c \in \mathbb{R}$,(1) $\mathbb{E}_w[c + Z] = c + \mathbb{E}_w[Z]$, and (2) $\mathbb{E}_w[cZ] = c\mathbb{E}_w[Z]$.*

Despite Lemma 1, and unlike risk-neutral expectation, the weighted expectation is *not* a linear operator on random variables, as the following example demonstrates.

Example 3. *Let Z_1 be distributed uniformly on $\{0, 1\}$ and let Z_2 be independently distributed uniformly on $\{0, 2\}$. Let $w(x) = x^2$. Then $\mathbb{E}_w[Z_1] = 1 \cdot w(\frac{1}{2}) = \frac{1}{4}$. Similarly, $\mathbb{E}_w[Z_2] = \frac{1}{2}$, so that $\mathbb{E}_w[Z_1] + \mathbb{E}_w[Z_2] = \frac{3}{4}$. On the other hand, because Z_1 and Z_2 are independent, the random variable $Z_1 + Z_2$ is uniform on $\{0, 1, 2, 3\}$, with weighted expectation*

$$\mathbb{E}_w[Z_1 + Z_2] = 1 \cdot (w(\frac{3}{4}) - w(\frac{1}{2})) + 2 \cdot (w(\frac{1}{2}) - w(\frac{1}{4})) + 3 \cdot w(\frac{1}{4}) = \frac{7}{8}.$$

[3] Rostek [30] studies in depth the preference model, termed "quantile maximization", implied by such weighting functions.

Mechanism Design. Back to mechanism design, any mechanism can be described by the allocation it makes and the payment it charges as a function of the buyer's report. For a report $v = (v_1, \ldots, v_m)$, we denote by $X(v)$ the random variable for the allocation, giving a probability to each possible allocation of the items in $\{0,1\}^m$. Similarly, $P(v)$ is the random variable for the payment when the report is v. $X(v)$ and $P(v)$ may be correlated. Importantly, common practices from mechanism design in the risk-neutral setting, like treating the allocation as a vector in $[0,1]^m$ or the payment as a real number (i.e. replacing the random variable of the payment with its expectation), are *with* loss of generality in our setting.

We assume that the buyer has additive utility for the items and is quasilinear with respect to payments: if she receives a set of items S for a payment p, her total value for this outcome is $\sum_{i \in S} v_i - p$. The buyer's weighted expected utility from the mechanism's outcome is $\mathbb{E}_w[v \cdot X(v) - P(v)]$; we say that a mechanism is incentive compatible (IC) for a buyer with weighting function w if for all possible values v, v' of the buyer, it holds that $\mathbb{E}_w[v \cdot X(v) - P(v)] \geq \mathbb{E}_w[v \cdot X(v') - P(v')]$. It is without loss of generality to express an incentive compatible mechanism in the form of a menu $\mathcal{M}$, with each menu item corresponding to a particular (allocation, payment) pair of correlated random variables (X, P). Then, the allocation and payment of a buyer with value v and weighting function w is given by the utility-maximizing menu item[4] $(X_w(v), P_w(v)) = \arg\max_{(X,P) \in \mathcal{M}} \mathbb{E}_w[v \cdot X - P]$. The revenue of the mechanism is given by $\text{Rev}_{\mathcal{M}}(w, \mathcal{D}) = \mathbb{E}[P(v)]$, where the expectation is with respect to the random valuation v (drawn from $\mathcal{D}$), as well as the random outcome of the payment random variable $P(v)$. A mechanism is individually rational (IR) if the buyer has non-negative expected utility when participating. Throughout the paper we focus on IC and IR mechanisms.

We slightly overload notation and let $\text{Rev}(w, \mathcal{D})$ denote the optimal revenue achievable by an incentive compatible mechanism from selling m items to a buyer with weighting function w and values drawn from $\mathcal{D}$. We will frequently drop w to indicate the risk-neutral optimal revenue, i.e. we use $\text{Rev}(\mathcal{D})$ to mean $\text{Rev}(\mathcal{I}, \mathcal{D})$ (recall that $\mathcal{I}$ is the risk-neutral weighting function, $\mathcal{I}(x) = x$), and $\text{DRev}(\mathcal{D})$ for the optimal revenue from a deterministic mechanism. Note that $\text{DRev}(w, \mathcal{D}) = \text{DRev}(w', \mathcal{D})$, for all w, w'.

In this paper we show that the best of $\text{SRev}(\mathcal{D})$ (or just SRev), the auction that sells each item separately at its optimal posted price, and $\text{BRev}(\mathcal{D})$ (or just BRev), the auction that sells the grand bundle as a single item, is a risk-robust approximation for a prospect theoretic buyer. For a risk-neutral buyer, the following result is known.

Theorem 1 ([3, 7])**.** *For a single, risk-neutral, additive bidder and any independent item distribution $\mathcal{D}$ it holds that*

$$\text{Rev}(\mathcal{I}, \mathcal{D}) \leq 2\text{BRev}(\mathcal{D}) + 4\text{SRev}(\mathcal{D}) \leq 6 \max\{\text{SRev}(\mathcal{D}), \text{BRev}(\mathcal{D})\}.$$

Quantifying Sensitivity to Risk. Ideally, we would like to give simple auctions that perform well for all weighting functions w, with respect to the optimal auction *tailored* for w. Unfortunately, as we see in the next section, for some of the mechanism classes we study such a goal is too optimistic without any restrictions on w. For instance, the $w(x)$ could take the value 1 for all x except $x = 0$. Here, the buyer's extreme optimism yields utility equal to that in her favorite outcome. Therefore, slight randomization in the outcomes (say with probability $\epsilon > 0$ the buyer pays nothing, but otherwise pays a very high price) would result in the buyer always having non-negative utility, making her an easy target for extracting arbitrary revenue. For some of our results we will therefore impose a mild restriction on w in order to escape these extreme situations, and otherwise make no assumptions (such as convexity or Lipschitzness).

[4]We assume that any ties are broken in favor of menu items with a higher expected price.

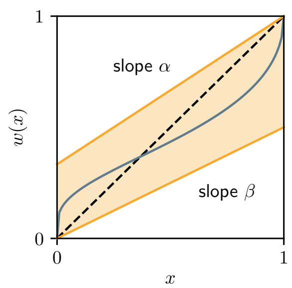

Figure 1: Any (α, β)-limited weighting function lies between the two red lines. The upper line has slope α and the lower has slope β.

Definition 2. *A weighting function w is* (α, β)-limited *if it satisfies (1) $w(x) \le \alpha(x-1)+1$, and (2) $w(x) \ge \beta x$.*

Geometrically, an (α, β)-limited weighting function w lies below the line with slope α passing through $(1, 1)$, and above the line with slope β passing through the origin; see Figure 1. The purpose of this definition is to control the slope of the weighting function as it approaches 0 and 1. Note that meaningful values for α and β lie in the range $[0, 1]$. As α and β approach 1, the buyer becomes less sensitive to risk; the risk-neutral weighting function is the unique $(1, 1)$-limited weighting function. For any $\alpha < \alpha'$ and any β, an (α, β)-limited weighting function is also (α', β)-limited. We stress that we don't use this restriction in all our results, and when we do use it, we need only one of the two sides of the bound, i.e. we ask for either $(\alpha, 0)$- or $(0, \beta)$-limited weighting functions.

3 Limits of the Model

In this section, we demonstrate how our model, absent any additional assumptions on the mechanism or the weighting function, can lead to absurd results. Such results were known before our work. Azevedo and Gottlieb [2] show that under assumptions on the weighting functions a principal can extract unbounded revenue from a CPT agent, simply by offering a bet on a single coin-flip. Furthermore, Ebert and Strack [15] show that CPT behavior gives rise to time inconsistency, allowing a seller to extract the buyer's entire wealth over multiple rounds of interaction. We reproduce similar results in our context for completeness and to illustrate the variety of behaviors possible in this model. In later sections, we develop restrictions on the mechanism which preclude this sort of unreasonable behavior. First, the following simple lemma is instructive.

Lemma 2. *For every distribution $\mathcal{D}$, constant $R \in \mathbb{R}_{\ge 0}$, and weighting function w such that there exists $x^* < 1$ with $w(x^*) = 1$, there exists a mechanism $\mathcal{M}$ such that $\text{Rev}_{\mathcal{M}}(w, \mathcal{D}) = R$.*

Proof. Consider the following lottery, where (positive) Z represents a transfer *to* the agent.

$$Z = \begin{cases} 0 & \text{with probability } x^* \\ \frac{-R}{1-x^*} & \text{with probability } 1 - x^*. \end{cases} \tag{2}$$

The agent's utility is $\mathbb{E}_w[Z] = \frac{-R}{1-x^*}(1 - w(x^*)) = 0$, while the seller's revenue is $\mathbb{E}[-Z] = \frac{R}{1-x^*}(1 - x^*) = R$. This lottery can be transformed into a mechanism for selling any number of items, by giving everything for free to the buyer, requiring only that she participates in the lottery. □

Lemma 2 relies on the dubious assumption that the buyer would assign no weight at all to an extremely negative—albeit potentially highly unlikely—outcome. However, even seemingly reasonable weighting functions can be exploited, as our next result shows.

Lemma 3. *For every distribution $\mathcal{D}$, constant $R \in \mathbb{R}_{\geq 0}$, and weighting function w such that there exists x^* with $1 > w(x^*) > x^*$, there exists a mechanism $\mathcal{M}$ such that $\text{REV}_{\mathcal{M}}(w, \mathcal{D}) = R$.*

Proof. Consider the following lottery, where (positive) Z represents a transfer to the agent.

$$Z = \begin{cases} a & \text{with probability } x^* \\ -\rho a & \text{with probability } 1 - x^*, \end{cases} \tag{3}$$

where $a > 0$. The expected value of an agent with weighting function w is $\mathbb{E}_w[Z] = aw(x^*) - \rho a\,(1 - w(x^*))$. Pick $\rho = \frac{w(x^*)}{1-w(x^*)}$; then, for all a, $\mathbb{E}_w[Z] = 0$. That is, the buyer has utility exactly zero for this lottery.

On the other hand, the expected revenue of the seller, who pays a with probability x^* and gets paid ρa with probability $1 - x^*$, is equal to

$$\mathbb{E}[-Z] = \rho a(1 - x^*) - ax^* = a \cdot \left(\frac{w(x^*)(1 - x^*)}{1 - w(x^*)} - x^*\right) = a \cdot \frac{w(x^*) - x^*}{1 - w(x^*)}.$$

The lemma follows by setting $a = R\frac{1-w(x^*)}{w(x^*)-x^*}$; similarly to Lemma 2, this lottery can be turned into an auction by giving all the items for free to the agent after participating in the lottery. □

We note that the conditions of Lemma 3 are satisfied for nearly all weighting functions implied by experiments in the literature; we refer the reader to [33, 34] for concrete examples. Furthermore, the issue exhibited by Lemma 3 persists even if one enforces ex-post individual rationality, so long as the seller is allowed to utilize a multi-round protocol.

Lemma 4. *For every distribution $\mathcal{D}$, constant $\epsilon > 0$ and weighting function w such that there exists x^* with $1 > w(x^*) > x^* + \frac{\epsilon}{1+\epsilon}$, there exists a multi-round, ex-post individually rational mechanism $\mathcal{M}$ such that $\text{REV}_{\mathcal{M}}(w, \mathcal{D}) = \mathbb{E}[\mathcal{D}]$.*

Proof. For simplicity we only prove the $m = 1$ item case; the general case is identical. Consider again the random transfer defined in (3). Picking $\rho = \frac{w(x^*)}{1-w(x^*)} - \epsilon$ provides the buyer strictly positive utility. The seller's revenue is equal to $\mathbb{E}[-Z] = a \cdot (\frac{w(x^*)(1-x^*)}{1-w(x^*)} - \epsilon(1 - x^*) - x^*)$, which is again strictly positive for every $a > 0$. By picking a and x^* appropriately the seller can thus make both $\mathbb{E}_w[Z]$ and $\mathbb{E}[-Z]$ very small positive numbers. This suffices to extract full buyer welfare as follows.

The buyer and seller will interact over T rounds. In the first round, the buyer reports a bid b. In rounds $t > 1$, the seller will offer lottery Z (and the buyer has the option to not participate), unless the seller has already extracted an amount larger than the bid b. After T rounds have passed, the item will be awarded to the buyer for free. Of course, since $\mathbb{E}_w[Z] > 0$, the buyer always chooses to participate in round t, and (in expectation) loses a little bit of money. By picking T large enough, the buyer eventually goes bankrupt at some intermediate round, but since she eventually gets the item this mechanism is in fact ex-post IR. Notice that this mechanism is also truthful! Precisely because when the buyer is calculating (in the first round) her expected utility from reporting b she thinks that she will "come out on top", and therefore is indifferent between all bids b (and thus reports her true value v). □

As the previous lemmas exhibit, practical mechanisms cannot hope to compete against the theoretically optimal revenue maximizing mechanism in this model, and thus this theory does not give accurate predictions for the simple mechanisms that we observe in practice. There are multiple ways to proceed. A natural one is to put restrictions on the weighting functions considered. Indeed, this is the approach taken by Chawla et al. [11] for the single item case, where the weighting function is restricted to be convex (therefore the buyer is always risk-averse). Another is to put restrictions on the mechanisms considered. In this paper we restrict our attention to specific mechanism classes; for some of our results this does not suffice and some mild restrictions on w are necessary as well.

4 Mechanism Classes

We define four classes of mechanisms. Recall that $\text{Rev}_{\mathcal{M}}(w, \mathcal{D})$ denotes the seller's expected revenue from a mechanism $\mathcal{M}$, given that the buyer has weighting function w and her values are distributed according to $\mathcal{D}$. We denote the expected revenue of the optimal mechanism in a class $\mathcal{C}$ by $\text{Rev}(w, \mathcal{D}, \mathcal{C})$. That is, $\text{Rev}(w, \mathcal{D}, \mathcal{C}) = \max_{\mathcal{M} \in \mathcal{C}} \text{Rev}_{\mathcal{M}}(w, \mathcal{D})$.

4.1 The class $\mathcal{C}_{dp}$ of deterministic price allocations.

First, we consider mechanisms which use randomness only in the allocation. That is, the seller offers a menu of distributions over the items, each at a fixed price. The buyer may pay the price for a distribution over the items, after which she receives a draw from the distribution. We call this class *deterministic price (DP)* mechanisms, and denote it by $\mathcal{C}_{dp}$. It will be convenient to think of a mechanism $\mathcal{M}$ in this class as a menu, where the buyer selects her favorite menu item, of the form (p, X), where p is the payment and X is a (possibly correlated) distribution over items. Observe that this class remains completely general for risk-neutral buyers.

Unfortunately, general deterministic price mechanisms are technically difficult to work with. The arbitrary correlation allowed between items (in the allocation) makes arguing about the buyer's expected utility problematic. Specifically, different buyer types order outcomes of X differently, and therefore could have wildly different expected weighted utility for the same distribution X (since arbitrary correlation allows us to assign arbitrary probabilities to outcomes); this property can be used to tailor to each type v an allocation $X(v)$ that is attractive only to this type. Our understanding of general $\mathcal{C}_{dp}$ mechanisms is therefore partial. We show that $\max\{\text{SRev}, \text{BRev}\}$ gives a doubly exponential (in the number of items) approximation to the optimal deterministic price mechanism. This trivially implies a constant approximation for a constant number of items; we leave it as an open problem whether a constant approximation can be achieved for an arbitrary number of items.

To mitigate the problems caused by arbitrary correlation, we also consider a special case of deterministic price mechanisms, which imposes a specific form of correlation on the distribution over allocations: we ask that the allocations in the support of the allocation distribution form a nested set. We term this class *nested deterministic price (NDP)* mechanisms and denote it by $\mathcal{C}_{ndp}$. We say a random variable X supported in $2^{[m]}$ is a *monotone lottery* if X is supported on a chain of subsets $S_1, \cdots, S_k$, $k \leq m$, such that $S_i \subset S_{i+1}$ for all $i \in [k-1]$. We use $\Delta_{\text{N}}(2^{[m]})$ to denote the set of such correlated distributions over the set of m items. For a mechanism $\mathcal{M} \in \mathcal{C}_{ndp}$ the allocation distributions for each menu item are restricted to be in $\Delta_{\text{N}}(2^{[m]})$. Observe that nested deterministic price mechanisms are again completely general for risk-neutral buyers. This is so because the optimal mechanism for a risk-neutral buyer can be specified in terms of the marginal probabilities of allocation for each item. For any marginal probabilities, we can find a monotone lottery having the same marginal probabilities.

Observation 1. *For any distribution $\mathcal{D}$, the class $\mathcal{C}_{ndp}$ of nested deterministic price mechanisms contains an optimal mechanism for a risk-neutral buyer. That is,* $\textsc{Rev}(\mathcal{I}, \mathcal{D}) = \textsc{Rev}(\mathcal{I}, \mathcal{D}, \mathcal{C}_{ndp})$.

Our main result for nested deterministic price mechanisms is that the seller cannot exploit the buyer's risk attitude at all in this class: $\textsc{Rev}(w, \mathcal{D}, \mathcal{C}_{ndp})$ is equal to $\textsc{Rev}(\mathcal{I}, \mathcal{D})$! This trivially implies that $\max\{\textsc{SRev}, \textsc{BRev}\}$ is a constant approximation to $\textsc{Rev}(w, \mathcal{D}, \mathcal{C}_{ndp})$ by Theorem 1.

4.2 The class $\mathcal{C}_{da}$ of deterministic-allocation mechanisms.

In Section 3, we showed that the seller can use randomized payments, utilizing very large positive transfers, to extract unbounded revenue from an optimistic buyer. But what power do randomized payments have when the buyer's willingness to participate is exogenous to the seller? Here, we consider mechanisms which satisfy two conditions: first, they sell only (deterministic) bundles of items, and second they cannot offer positive transfers, that is, monetary transfers from the seller to the buyer. We refer to these as *deterministic allocation* (DA) mechanisms, and denote it by $\mathcal{C}_{da}$. Formally, every $\mathcal{M} \in \mathcal{C}_{da}$ offers one lottery for each subset of items $S \subseteq [m]$. The lottery for a subset S always allocates S but charges a random (non-negative) payment.

Observe that this class of mechanisms is not fully general for a risk-neutral buyer, but, since allocations are deterministic, DA mechanisms are equivalent to fully deterministic mechanisms for a risk-neutral buyer. We observe that a seller with knowledge of w can use a DA mechanism to extract more revenue from a RDUT buyer than would be possible to extract from a risk-neutral buyer via any mechanism, i.e. $\textsc{Rev}(w, \mathcal{D}, \mathcal{C}_{da}) > \textsc{Rev}(\mathcal{I}, \mathcal{D})$. In fact, the gap between the two is unbounded. This is possible because an optimistic buyer is willing to gamble that she will not have to make a payment. Our main positive result for DA mechanisms is that, if the weighting function is $(\alpha, 0)$-limited, $\textsc{DRev}(\mathcal{I}, \mathcal{D})$ is an α approximation to $\textsc{Rev}(w, \mathcal{D}, \mathcal{C}_{da})$; therefore $\max\{\textsc{SRev}, \textsc{BRev}\}$ is a 6α approximation to $\textsc{Rev}(w, \mathcal{D}, \mathcal{C}_{da})$, for all distribution $\mathcal{D}$. As a special case of this result, we get that the seller can use randomized payments to extract extra revenue only from optimistic buyers; if $w(x) \le x$ for all x, i.e. $\alpha = 1$, there is no loss in the approximation.

4.3 The class $\mathcal{C}_b$ of binary-lottery mechanisms.

Finally, we consider a generalization of the mechanism format studied by Chawla et al. [11]. They showed that a menu of lotteries supported on only two outcomes—receive the item and pay, or pay nothing and receive nothing—were sufficient to extract the optimal revenue if there is only a single item for sale and the buyer has a convex weighting function. We consider an extension to multiple items which offers, for each *subset* of items, a (potentially uncountable) menu of binary lotteries. Let $\mathcal{C}_b$ denote the following class of auctions. Each $\mathcal{M} \in \mathcal{C}_b$ contains binary lotteries for subsets $S \subseteq [m]$. A lottery for a subset S is of the following format:

$$(X_S, P_S) = \begin{cases} (S, p_S) & \text{with probability } q_S \\ (\emptyset, 0) & \text{with probability } 1 - q_S. \end{cases}$$

That is, either get the subset S and pay p_S (with probability q_S), or get nothing and pay nothing (with probability $1 - q_S$). Note that this does constitute a significant restriction of the design space. However, this format is still quite flexible: considered as a direct-revelation mechanism, for each type we specify a subset of the items, a probability of allocation, and a payment. Our main positive result for binary-lottery mechanisms is that, if the weighting function is $(0, \beta)$-limited, $\textsc{Rev}(\mathcal{I}, \mathcal{D})$ is a β approximation to $\textsc{Rev}(w, \mathcal{D}, \mathcal{C}_b)$; combined with Theorem 1 this implies a 6β approximation for $\max\{\textsc{SRev}, \textsc{BRev}\}$.

5 Deterministic Price Mechanisms

We first investigate general deterministic price mechanisms. We show that the optimal revenue of a deterministic price mechanism on independent items for a RDUT buyer can be upper bounded by doubly exponential times the optimal *risk-neutral* revenue of some items and the welfare on the distribution of the remaining items.

Theorem 2. *Let w be a weighting function, $\mathcal{D}_1$ be an independent distribution over m_1 items, and $\mathcal{D}_2$ be an independent distribution over m_2 items. Let $\mathcal{D} = \mathcal{D}_1 \times \mathcal{D}_2$ be an independent distribution over m items, where $m = m_1 + m_2$. Then, for a single, additive bidder it holds that*

$$\textsc{Rev}(w, \mathcal{D}, \mathcal{C}_{dp}) \leq 2^{2^{m_1(m - \frac{1}{2}\log m_1)}} \textsc{Rev}(\mathcal{I}, \mathcal{D}_1) + \mathrm{VAL}\left(\mathcal{D}_2\right) \quad .$$

Using standard techniques we get the following corollary.

Corollary 1. *For a single, additive bidder and any independent item distribution $\mathcal{D}$, it holds that $\textsc{Rev}(w, \mathcal{D}, \mathcal{C}_{dp}) \in O(2^{m2^m}) \max\{\textsc{SRev}, \textsc{BRev}\}$.*

Proof. Let T_1 and T_2 be the random variables of the sum of value of items in $\mathcal{D}_1$ and $\mathcal{D}_2$ respectively. We use the notation $\mathbb{1}_E$ for the indicator variable of an event E. Notice that $\textsc{Rev}_{\mathcal{M}}(w, \mathcal{D}, C_{dp}) = \mathbb{E}[p(v_1, v_2) \cdot \mathbb{1}_{T_1 \geq T_2}] + \mathbb{E}[p(v_1, v_2) \cdot \mathbb{1}_{T_1 < T_2}]$. Each of these terms can be upper bounded using Theorem 2. For example, for the first term we have:

$$\mathbb{E}[p(v_1, v_2) \cdot \mathbb{1}_{T_1 \geq T_2}] \leq 2^{2^{m_1(m - \frac{1}{2}\log m_1)}} \textsc{Rev}(\mathcal{I}, \mathcal{D}_1) + \mathrm{VAL}\left(\mathcal{D}_2 \cdot \mathbb{1}_{T_1 \geq T_2}\right),$$

Notice that by posting a price of $\mathrm{VAL}\left(\mathcal{D}_2\right) = \mathbb{E}[T_2]$ for the grand bundle (of m_1 items) in $\mathcal{D}_1$, we can make revenue $\mathbb{E}[T_2 \cdot \mathbb{1}_{T_1 \geq T_2}]$. Therefore $\mathbb{E}[T_2 \cdot \mathbb{1}_{T_1 \geq T_2}] \leq \textsc{BRev}(\mathcal{D}_1)$. Thus

$$\mathbb{E}_{v_1 \sim \mathcal{D}_1}[p(v_1, v_2) \cdot \mathbb{1}_{T_1 \geq T_2}] \leq 2^{2^{m_1(m - \frac{1}{2}\log m_1)}} \textsc{Rev}(\mathcal{I}, \mathcal{D}_1) + \textsc{BRev}(\mathcal{D}_1).$$

The symmetric bound holds for $\mathbb{E}[p(v_1, v_2) \cdot \mathbb{1}_{T_1 < T_2}]$. Adding the two inequalities and combining with the upper bound on the optimal risk-neutral revenue (Theorem 1) we get the corollary. □

Though this approximation is doubly exponential in the number of items, we do get a constant approximation when the number of items is a constant. Notably, for the case of two items, we get $\textsc{Rev}(w, \mathcal{D}, \mathcal{C}_{dp}) \leq 17\textsc{SRev}$; an improved analysis can reduce this to a factor of 2. We leave it as an open problem whether a constant approximation is possible for an arbitrary number of items.

Before we proceed with the proof of Theorem 2 we briefly comment on the technical obstacles that lead to this approximation factor. As we've mentioned in the introduction, for a risk-neutral buyer, statements similar to Theorem 2 are known, for example $\textsc{Rev}(\mathcal{D}_1 \times \mathcal{D}_2) \leq \textsc{Rev}(\mathcal{D}_1) + \mathrm{VAL}(\mathcal{D}_2)$ ([3, 22]). The proof of this statement is roughly as follows. For every $v_2 \in \mathcal{D}_2$ construct an auction $\mathcal{M}^{v_2}$ for $\mathcal{D}_1$. When the buyer reports some $v_1 \in \mathcal{D}_1$, $\mathcal{M}^{v_2}$ copies the allocation (for the items in $\mathcal{D}_1$) and payment rule of the optimal auction, $\mathcal{M}$, for $\mathcal{D}_1 \times \mathcal{D}_2$, and then slightly adjusts the payments to cover for the lost value (i.e. expected value from the items in $\mathcal{D}_2$ that should not be paid in $\mathcal{M}^{v_2}$). The result of this adjustment is that the utility of reporting v_1 in $\mathcal{M}^{v_2}$ is equal to the utility of reporting (v_1, v_2) in $\mathcal{M}$, and therefore $\mathcal{M}^{v_2}$ is incentive compatible. Then, the optimal revenue for $\mathcal{D}_1$ is at least the revenue $\mathcal{M}^{v_2}$, which is equal to $\textsc{Rev}(\mathcal{D}_1 \times \mathcal{D}_2)$ minus the reimbursement (which is at most $\mathrm{VAL}(\mathcal{D}_2)$). Unfortunately, such adjustments to the payment do not work here. Replacing the (random) allocation of items in $\mathcal{D}_2$ with a deterministic payment gives different weighted expected value to different types, since different types order outcomes differently. For example, a buyer with type $t = (v_1, v_2, v_3)$ such that $v_3 > v_1 + v_2$ prefers outcome

$\{3\}$ to outcome $\{1,2\}$; the calculation of the weighted expected utility depends on the order over outcomes, and therefore, because of the IC constraints between types with different orders, we can't naively replace random outcomes with deterministic payments (or, more accurately, deterministic reimbursements).

Our approximation factor depends on the number of all possible *valid* orderings over outcomes. By valid ordering we mean that there exist types who order outcomes of X in this ordering (for example, an ordering where outcome $\{1\}$ is preferred to $\{1,2\}$ is not valid) . A loose upper bound on the number of valid orderings is $(2^m)!$; in our proof we're able to do much better. A couple of new additional ideas are necessary in order to deal with the subtleties of our setting, but otherwise our proof follows a similar structure to the one described above.

Proof of Theorem 2. Let $\mathcal{M}$ be the revenue optimal auction (in auction class C_{dp}) for a buyer with weighting function w and value distribution $\mathcal{D} = \mathcal{D}_1 \times \mathcal{D}_2$; let $X(v)$ and $p(v)$ be the allocation and payment rule of $\mathcal{M}$. We use the notation $\mathbb{1}_E$ for the indicator variable of an event E. For ease of notation we will write $i \in [m_j]$ instead of "item i in the support of $\mathcal{D}_j$", and $v_i(j)$ for the value an item j in the type $v_i \in \mathcal{D}_i$[5]. We prove the theorem by giving an auction for selling m_1 items to a *risk-neutral* buyer whose values are drawn from $\mathcal{D}_1$. Specifically, first sample v_2 from $\mathcal{D}_2$; we construct a deterministic-price mechanism $\mathcal{M}^{v_2} = (X^{v_2}, p^{v_2})$ that is incentive compatible for a risk-neutral buyer, and then lower bound its revenue.

We would like to relate the revenue $\mathcal{M}^{v_2}$ with the revenue of $\mathcal{M}$. As we've already discussed though, the natural choice of copying the allocation for items $i \in [m_1]$ from $\mathcal{M}$ and then adjusting the payments does not work, because the resulting auction is not incentive compatible. In this proof we will restrict our attention to a certain subset R^{v_2} of the support of $\mathcal{D}_1$, such that all types in R^{v_2} order outcomes the same way. We then copy the decisions of $\mathcal{M}$ only in that region; when $v_1 \notin R^{v_2}$ the buyer is allocated her favorite menu item (from the ones designed for types in R^{v_2}).

Use $S_1, \dots, S_{2^m}$ to denote the 2^m distinct subsets of m items. A necessary condition for an ordering over items to be valid is inclusion partial order: if $S_i \subset S_j$, then S_i is before S_j in the order. Brightwell and Tetali [4] have shown that there are at most $2^{2^{m_1}(m_1 - \frac{1}{2}\log m_1 - \log e)}$ ways to order the subsets of m_1 items such that the orders satisfy the inclusion partial order property. For each of those ordering, there are at most $\binom{2^m - 2^{m_2}}{2^{m_1}}/(2^{m_2}!)$ [6] different ways to combine an ordering on subsets of m_1 items and an ordering on subsets of m_2 items (that satisfy the inclusion partial order property) to obtain an ordering on the subsets of their union that satisfies the inclusion partial order property. Thus, there are at most $N = 2^{2^{m_1}(m_1 - \frac{1}{2}\log m_1 - \log e)} \cdot \binom{2^m}{2^{m_1}}/(2^{m_2}!)$ different orderings of $S_1, \dots, S_{2^m}$ when v_2 is fixed. Therefore we can partition the support of $\mathcal{D}_1$ into regions $R_1, \dots, R_N$, such that in each region R_i, all types have the same preference ordering over subsets of items.

Let R^{v_2} denote the region from which $\mathcal{M}$ extracts the most revenue (weighted by the probability of being in the region). Assume wlog that the order of subsets in this region is $S_1 \leq S_2 \leq \cdots \leq S_{2^m}$. For every $v_1 \in R^{v_2}$, we want the risk-neutral buyer's utility of reporting v_1 in $\mathcal{M}^{v_2}$ to be the same as the RDUT buyer's utility of reporting (v_1, v_2) in $\mathcal{M}$. The latter term is equal to

$$u_w((v_1, v_2), \mathcal{M}) = -p(v_1, v_2) + \sum_{i \in [m_1]} a_i v_1(i) + \sum_{j \in [m_2]} b_j v_2(j) \ , \tag{4}$$

for some non-negative constants $a_1, \dots, a_{m_1}$ and $b_1, \dots, b_{m_2}$. Specifically, think of $X(v)$ as a vector

[5] $v_i \in \mathcal{D}_i$ is an m_i dimensional vector.

[6] $\binom{2^m}{2^{m_1}}$ counts the number of complete orders with the order of all subsets of $[m_1]$ items fixed. Dividing by $(2^{m_2}!)$ removes the orders in which the subsets of $[m_2]$ items are ordered incorrectly.

in $\mathbb{R}^{2^m}_{\geq 0}$ where the i-th entry, denoted by $X(v)[S_i]$, is the probability that the buyer gets exactly the set S_i. Then $a_i = \sum_{k\in[2^m],S_k\ni i}\left(w(\sum_{j\geq k}^{2^m} X(v_1,v_2)[S_j]) - w(\sum_{j>k}^{2^m} X(v_1,v_2)[S_j])\right)$.

Each term in the sum is nonnegative (since the weighting function is non-decreasing) and the sum of a_is telescopes, that is $\sum_{k\in[2^m]}\left(w(\sum_{j\geq k}^{2^m} X(v_1,v_2)[S_j]) - w(\sum_{j>k}^{2^m} X(v_1,v_2)[S_j])\right) = 1$. Therefore we can conclude that all coefficients a_i are in $[0,1]$, and since they add up to 1, we can think of them as probabilities. Similarly, for the b_js.

Since $\mathcal{M}^{v_2}$ is designed for a risk-neutral buyer, we can think of $X^{v_2}(v_1)$ as a vector in $\mathbb{R}^{m_1}_{\geq 0}$. The i-th entry of $X^{v_2}(v_1)$ is the marginal probability that the buyer gets item $i \in [m_1]$. Now, for every $v_1 \in R^{v_2}$, we set the allocation to be $X^{v_2}(v_1) = \begin{bmatrix} a_1 & \dots & a_{m_1}\end{bmatrix}^T$, and the payment to be $p^{v_2}(v_1) = p(v_1,v_2) - \sum_{j\in[m_2]} b_j v_2(j)$. It is immediate that the utility of a risk-neutral buyer when reporting $v_1 \in R^{v_2}$ is exactly the RHS of (4). Since $\mathcal{M}$ is IC, we get that, in $\mathcal{M}^{v_2}$, if the (risk-neutral) buyer's type is in R^{v_2}, then she has no incentive to report a different type in R^{v_2}. Finally, for every $v_1 \notin R^{v_2}$, set $(X^{v_2}(v_1), p^{v_2}(v_1))$ to be the best menu item among the ones already added to $\mathcal{M}^{v_2}$ for types in R^{v_2}. Therefore, $\mathcal{M}^{v_2}$ is trivially IC for $v_1 \notin R^{v_2}$ as well. Intuitively, seen as a menu, our overall auction $\mathcal{M}^{v_2}$ copies (and adjusts) the menu of $\mathcal{M}$ for types (v_1, v_2) where $v_1 \in R^{v_2}$.

It remains to bound the revenue of $\mathcal{M}^{v_2}$. The revenue of $\mathcal{M}^{v_2}$ is at least its revenue in R^{v_2}:

$$\textsc{Rev}_{\mathcal{M}^{v_2}}(\mathcal{I},\mathcal{D}_1) \geq \mathbb{E}_{v_1\sim\mathcal{D}_1}\left[\left(p(v_1,v_2) - \sum_{j\in[m_2]} b_j v_2(j)\right)\cdot \mathbb{1}_{v_1\in R^{v_2}}\middle| v_2\right].$$

Since the revenue of $\mathcal{M}^{v_2}$ cannot exceed the optimal revenue, we have

$$\begin{aligned}
\mathbb{E}_{v_1\sim\mathcal{D}_1}\left[\left(p(v_1,v_2) - \sum_{j\in[m_2]} b_j v_2(j)\right)\cdot \mathbb{1}_{v_1\in R^{v_2}}\middle| v_2\right] &\leq \textsc{Rev}(\mathcal{I},\mathcal{D}_1)\\
\mathbb{E}_{v_1\sim\mathcal{D}_1}\left[\left(p(v_1,v_2) - \sum_{j\in[m_2]} b_j v_2(j)\right)\middle| v_2\right] &\leq N\cdot\textsc{Rev}(\mathcal{I},\mathcal{D}_1)\\
\mathbb{E}_{v_1\sim\mathcal{D}_1}\left[p(v_1,v_2)|v_2\right] &\leq N\cdot\textsc{Rev}(\mathcal{I},\mathcal{D}_1) + \mathbb{E}\left[\left(\sum_{j\in[m_2]} b_j v_2(j)\right)|v_2\right]\\
\mathbb{E}_{v_1\sim\mathcal{D}_1}\left[p(v_1,v_2)|v_2\right] &\leq N\cdot\textsc{Rev}(\mathcal{I},\mathcal{D}_1) + \text{VAL}\left(\mathcal{D}_2|v_2\right).
\end{aligned}$$

Finally, sum this inequality across the support of $\mathcal{D}_2$ to get

$$\begin{aligned}
\sum_{v_2\sim\mathcal{D}_2} Pr[v_2]\cdot\mathbb{E}_{v_1\sim\mathcal{D}_1}[p(v_1,v_2)|v_2] &\leq \sum_{v_2\sim\mathcal{D}_2} Pr[v_2]\cdot\left(N\cdot\textsc{Rev}(\mathcal{D}_1) + \text{VAL}\left(\mathcal{D}_2|v_2\right)\right)\\
\mathbb{E}[p(v_1,v_2)] &\leq N\cdot\textsc{Rev}(\mathcal{I},\mathcal{D}_1) + \text{VAL}\left(\mathcal{D}_2\right).
\end{aligned}$$

Observing that $N = 2^{2^{m_1}(m_1-\frac{1}{2}\log m_1 - \log e)}\cdot\binom{2^m}{2^{m_1}}/(2^{m_2}!) < 2^{2^{m_1}(m-\frac{1}{2}\log m_1)}$ concludes the proof. $\square$

5.1 Nested Deterministic Price Mechanisms

Our main result is that the class of nested deterministic price mechanisms does not offer the seller any means of exploiting the buyer's risk attitude: the optimal revenue within the class is equivalent to the optimal revenue obtainable from a risk-neutral mechanism.

Theorem 3. *Let w be an invertible weighting function and $\mathcal{D}$ be an independent distribution supported in $\mathbb{R}^m_{\geq 0}$. Then, for a single, additive bidder it holds that* $\textsc{Rev}(w,\mathcal{D},\mathcal{C}_{ndp}) = \textsc{Rev}(\mathcal{I},\mathcal{D})$.

Combining this result with Theorem 1 of Babaioff et al. [3] we get the following corollary.

Corollary 2. *Let w be any invertible weighting function and $\mathcal{D}$ be any independent distribution supported in $\mathbb{R}^m_{\geq 0}$. Then, for a single, additive bidder it holds that*

$$\textsc{Rev}(w, \mathcal{D}, \mathcal{C}_{ndp}) \leq 6 \max\{\textsc{SRev}(\mathcal{D}), \textsc{BRev}(\mathcal{D})\}.$$

We prove Theorem 3 in two lemmas. We start by showing that for any invertible weighting function, there exists an NDP mechanism which recovers the optimal risk-neutral revenue.

Lemma 5. *Let w be an invertible weighting function and $\mathcal{D}$ be any distribution supported in $\mathbb{R}^m_{\geq 0}$. Then $\textsc{Rev}(w, \mathcal{D}, \mathcal{C}_{ndp}) \geq \textsc{Rev}(\mathcal{I}, \mathcal{D})$.*

Proof. Given a mechanism $\mathcal{M}$ such that $\textsc{Rev}_{\mathcal{M}}(\mathcal{I}, \mathcal{D}) = \textsc{Rev}(\mathcal{I}, \mathcal{D})$ (i.e. a revenue optimal mechanism for a risk-neutral buyer), let $X(v)$ and $p(v)$ be the allocation and payment rule (respectively) of $\mathcal{M}$. Since this buyer is risk-neutral, $p(v) \in \mathbb{R}_{\geq 0}$ is some deterministic payment, and we will assume that $X(v)$, the random variable for the allocation, is a monotone lottery (which is without loss of generality by Observation 1). We define a mechanism $\tilde{\mathcal{M}} = \big(\tilde{X}(v), \tilde{p}(v)\big)$ such that $\textsc{Rev}_{\tilde{\mathcal{M}}}(w, \mathcal{D}) = \textsc{Rev}_{\mathcal{M}}(\mathcal{I}, \mathcal{D})$.

Fix v. Let $S_1, \cdots, S_k$ be the support of $X(v)$, where $S_i \subset S_{i+1}$ for $i \in [k]$, and let $1 - F_i = \Pr[S_i \subseteq X(v)]$. Then a risk-neutral buyer's utility when participating in $\mathcal{M}$ is $u(v, \mathcal{M}) = \sum_{i=1}^{k} \big(v(S_i) - v(S_{i-1})\big)(1 - F_i)$, where we take $S_0 = \emptyset$. Let $1 - \tilde{F}_i = w^{-1}(1 - F_i)$, and define $\tilde{X}(v)$ such that $\Pr\Big[\tilde{X}(v) = S_i\Big] = \tilde{F}_{i+1} - \tilde{F}_i$. Also, let $\tilde{p}(v) = p(v)$. A risk-sensitive buyer with weighting function w and *any* valuation v' has weighted expected utility for the lottery $(\tilde{X}(v), \tilde{p}(v))$ equal to

$$\begin{aligned} u_w(v', \tilde{X}(v), \tilde{p}(v)) &= \sum_{i=1}^{k} \big(v'(S_i) - v'(S_{i-1})\big) w(1 - \tilde{F}_i) - \tilde{p}(v) \\ &= \sum_{i=1}^{k} \big(v'(S_i) - v'(S_{i-1})\big)(1 - F_i) - p(v) \\ &= u(v', X(v), p(v)) \end{aligned}$$

The first equality follows because $v'(S)$ is monotone in S and $S_i \subseteq S_{i+1}$. Because this equality holds for every valuation v', $\big(\tilde{X}(v), \tilde{p}(v)\big)$ is an IC and IR mechanism for a buyer with weighting function w, and furthermore obtains the same revenue from that buyer as $\mathcal{M}$ obtains from a risk-neutral buyer. □

Next, we show the converse: that we can construct a mechanism for a risk-neutral buyer which obtains the same revenue as any DP mechanism for a buyer with weighting function w.

Lemma 6. *Let w be any weighting function and $\mathcal{D}$ any distribution supported in $\mathbb{R}^m_{\geq 0}$. Then $\textsc{Rev}(w, \mathcal{D}, \mathcal{C}_{ndp}) \leq \textsc{Rev}(\mathcal{I}, \mathcal{D})$.*

Proof. Consider a mechanism $\mathcal{M} \in \mathcal{C}_{ndp}$. Let $X(v)$ and $p(v)$ be the allocation and payment rule, respectively, of $\mathcal{M}$, where $X(v)$ is a random variable in $\Delta_{\text{N}}(2^{[m]})$ and $p(v) \in \mathbb{R}_{\geq 0}$. We construct a mechanism $\tilde{\mathcal{M}} = \big(\tilde{X}(v), \tilde{p}(v)\big)$ for a risk-neutral buyer such that $\textsc{Rev}_{\tilde{\mathcal{M}}}(\mathcal{I}, \mathcal{D}) = \textsc{Rev}_{\mathcal{M}}(w, \mathcal{D})$.

Fix v. $X(v)$ is a monotone lottery by definition of $\mathcal{C}_{ndp}$, so let $S_1, \cdots, S_k$ be the support of $X(v)$, where $S_i \subset S_{i+1}$ for $i \in [k]$, and let $1 - F_i = \Pr[S_i \subseteq X(v)]$. Then the utility of an RDUT buyer is $u_w(v, X(v), p(v)) = \sum_{i=1}^{k} (v(S_i) - v(S_{i-1}))\, w(1 - F_i)$, where we take $S_0 = \emptyset$. Let $1 - \tilde{F}_i = w(1 - F_i)$, and define $\tilde{X}(v)$ such that $\Pr\left[\tilde{X}(v) = S_i\right] = \tilde{F}_{i+1} - \tilde{F}_i$. Lastly, let $\tilde{p}(v) = p(v)$. A risk-neutral buyer with *any* valuation v' has expected utility for the lottery $(\tilde{X}(v), \tilde{p}(v))$ equal to

$$\begin{aligned} u(v', \tilde{X}(v), \tilde{p}(v)) &= \sum_{i=1}^{k} \big(v'(S_i) - v'(S_{i-1})\big)(1 - \tilde{F}_i) - \tilde{p}(v) \\ &= \sum_{i=1}^{k} \big(v'(S_i) - v'(S_{i-1})\big) w(1 - F_i) - p(v), \end{aligned}$$

which is just $u_w(v', X(v), p(v))$. Because this equality holds for every valuation v', $\big(\tilde{X}(v), \tilde{p}(v)\big)$ is an IC, IR mechanism for a buyer with weighting function w, and furthermore obtains the same revenue from a buyer with weighting function w as $\mathcal{M}$ obtains from a risk-neutral buyer. □

Observe that the assumption of monotone lotteries was critical to the proof of Lemma 6. If $X(v)$ were an arbitrary distribution over subsets $S \in 2^{[m]}$, a buyer with valuation v' would order the outcomes differently from v. This would make it impossible to define the unweighted probability of allocation in the mechanism $\tilde{\mathcal{M}}$ in a way that would be simultaneously consistent with the weighted probability assigned to the outcome by all valuations v'.

Indeed a general deterministic-price mechanism (without the restriction to monotone lotteries) could exploit this discrepancy to obtain more revenue than a risk-neutral mechanism. That is, Lemma 6 does not hold for the class $\mathcal{C}_{dp}$. We show such an example below.

Claim 1. *There exist a distribution $\mathcal{D}$ over two items, and a weighting function w, such that $\textsc{Rev}(w, \mathcal{D}, \mathcal{C}_{dp}) > \textsc{Rev}(\mathcal{I}, \mathcal{D})$.*

Proof. Let $\mathcal{D}_1, \mathcal{D}_2$ be independent and identical uniform distributions on $\{1, 3\}$. The revenue optimal auction that sells the two items to a risk-neutral buyer is the deterministic auction that sells the bundle of two items at the price 4. So $\textsc{Rev}(\mathcal{I}, \mathcal{D}_1 \times \mathcal{D}_2) = 4 \times \frac{3}{4} = 3$. Consider the weighting function

$$w(p) = \begin{cases} 0, & p \le \frac{1}{2} \\ 4p - 2, & \frac{1}{2} < p < \frac{3}{4} \\ 1, & \frac{3}{4} \le p \end{cases}.$$

Consider the auction $\mathcal{M}$ selling the two items in the following way: if the buyer reports type $(1, 1)$ she gets the first item with probability $\frac{1}{2}$ and independently, gets the second item with probability $\frac{1}{2}$, and the buyer pays 1 to the seller. Otherwise, the buyer gets both items and pays 4. It is easy to see that $\mathcal{M}$ is incentive compatible for a buyer with weighting function w. Furthermore, $\textsc{Rev}_{\mathcal{M}}(w, \mathcal{D}_1 \times \mathcal{D}_2) = 1 \times \frac{1}{4} + 4 \times \frac{3}{4} = \frac{13}{4} > 3$. □

6 Deterministic Allocation Mechanisms

We now turn to deterministic allocation mechanisms which randomize the payment, but which are restricted to offer only deterministic bundles and charge only positive payments ex post. Unlike the examples of Section 3, these mechanisms cannot offer a positive transfer in order to induce the buyer to pay more. Instead, the value of the items received must induce the buyer to pay. How much revenue can the seller then obtain?

As the next example shows, with randomized payments the seller can obtain strictly more revenue from an RDUT buyer than from a risk-neutral buyer, even with one item and a mechanism that *always* allocates the item.

Example 4. *Let* $w(x) = x^{1/a}$ *for* $a \geq 1$*; note that a buyer with such a weighting function is strictly optimistic. Suppose there is one item available, and* $\mathcal{D} = U[0,1]$*. Consider the mechanism that offers the item for a randomized payment* P *which is* p *w.p.* $1-q$ *and* 0 *otherwise. For any value* v *for the item, the utility of the buyer is* $v - p(1-w(q))$*. In other words, the buyer purchases the item if* $v \geq p(1-w(q))$*. The total expected revenue is thus*

$$p(1-q)\left(1 - F\left(p(1-w(q))\right)\right) = p(1-q)\left(1 - p(1-w(q))\right). \tag{5}$$

Fix q*; we will solve for the optimal* p*. By the first derivative test, we find that the optimal price is* $p_q = \frac{1}{2(1-w(q))}$*. Substituting this into* (5)*, we have*

$$\textsc{Rev}(q) = \frac{1-q}{4\big(1-w(q)\big)} = \frac{1-q}{4\big(1-q^{1/a}\big)}.$$

By L'Hôpital's rule, $\lim_{q \to 1} \textsc{Rev}(q) = \frac{a}{4}$*. Observing that* $\textsc{Rev}(\mathcal{D}) = \max_x x \cdot Pr[v \geq x] = \frac{1}{4}$*, we see that the revenue approaches* $a \cdot \textsc{Rev}(\mathcal{D})$*.*

The main result of this section shows that Example 4 gives the largest possible gap between the revenue of a DA mechanism and the revenue obtainable via any deterministic mechanism from a risk-neutral buyer. Note that $x^{1/a}$ is an $(\frac{1}{a}, 1)$-limited weighting function. We show that no deterministic-allocation mechanism can increase the seller's revenue by a factor more than α^{-1} over the risk-neutral optimum for any $(\alpha, 0)$-limited weighting function. As a special case of this result, we get that the seller can use randomized payments only to extract extra revenue from optimistic buyers; for pessimistic buyers, the revenue extractable via DA mechanisms is equal to the revenue obtainable via fully deterministic mechanisms.

Theorem 4. *For a single, additive bidder, any* $(\alpha, 0)$*-limited weighting function* w *and any independent distribution* $\mathcal{D}$ *it holds that*

$$\textsc{Rev}(w, \mathcal{D}, \mathcal{C}_{da}) \leq \alpha^{-1} \text{D}\textsc{Rev}(\mathcal{D}).$$

Combining Theorems 4 and 1 we get the following corollary.

Corollary 3. *For a single, additive bidder, any* $(\alpha, 0)$*-limited weighting function* w *and any independent distribution* $\mathcal{D}$*, it holds that* $\textsc{Rev}(w, \mathcal{D}, \mathcal{C}_{da}) \leq 6\alpha^{-1} \max\{\text{S}\textsc{Rev}, \text{B}\textsc{Rev}\}$*.*

Before proving Theorem 4, we show that it is without loss of generality to assume that the payment variables in a DA mechanism have a very simple form. Namely, we show that it suffices to consider Bernoulli-distributed payments: with probability q the payment is some positive p, and otherwise the payment is zero.

Lemma 7. *Fix a distribution* $\mathcal{D}$ *and a continuous,* $(\alpha, 0)$*-limited weighting function* w *with* $\alpha > 0$*. For every deterministic-allocation mechanism* $\mathcal{M}$ *with payments* P_S *with discrete support in* $[0, \infty)$*, there exists a deterministic-allocation mechanism* $\mathcal{M}'$ *with binary payments* P'_S *such that*

$$\textsc{Rev}_{\mathcal{M}}(w, \mathcal{D}) \leq \textsc{Rev}_{\mathcal{M}'}(w, \mathcal{D}).$$

Furthermore, when $w(x) \leq x$ *for all* x*,* P'_S *is a deterministic price.*

Proof. For this proof, it will be convenient to work with the dual of the weighting function $w^\dagger(x) = 1 - w(1-x)$. The proof of the following claim can be found in Appendix C.

Claim 2. $\mathbb{E}_w[-X] = -\mathbb{E}_{w^\dagger}[X]$

Fix a set S with associated payment P_S in $\mathcal{M}$. $\mathcal{M}'$ has the same allocation as $\mathcal{M}$, but different payments P'_S, such that

$$\mathbb{E}_{w^\dagger}[P_S] = \mathbb{E}_{w^\dagger}\left[P'_S\right] \text{ and} \tag{6}$$

$$\mathbb{E}[P_S] \leq \mathbb{E}\left[P'_S\right], \tag{7}$$

that is, the expected payment to the seller (aka the revenue) is non-decreasing, while the (weighted) expectation of the buyer's payment is the same. To see why these equations imply the result, notice that for every set S and type v the utility of a buyer satisfies

$$u_w(v, S, P_S) = \mathbb{E}_w[v(S) - P_S] = v(S) - \mathbb{E}_{w^\dagger}[P_S] = v(S) - \mathbb{E}_{w^\dagger}\left[P'_S\right] = u(v, S, P'_S),$$

where the second equality follows from Lemma 1 (since the allocation is deterministic) and Claim 2, and the third equality follows from Equation (6). Therefore the utility of each menu item is the same, and thus for every type v the buyer will make the same selection in $\mathcal{M}'$ as in $\mathcal{M}$.

Now we show that we can find prices P'_S satisfying (6) and (7). For ease of notation, we consider a single set S and omit the subscripts. Let P be supported on $\{p_1, \cdots, p_k\}$, where $\Pr[P = p_i] = q_i$ for all $i \in [k]$. For some $p \in \mathbb{R}_{\geq 0}$ and $q \in [0, 1]$ to be specified, let $P' = p$ w.p. q and 0 w.p. $1 - q$. Let F_P be the CDF of P. Then the conditions become

$$pw^\dagger(q) = \int_0^\infty w^\dagger\big(1 - F_P(z)\big)dz \qquad \text{and} \qquad pq \geq \int_0^\infty \big(1 - F_P(z)\big)dz.$$

Let $q = \max_{i \in [k]} \frac{q_i}{w^\dagger(q_i)}$. Observe that if $w(x) \leq x$ for all x, we can take $q = 1$. Set $p = \frac{1}{w^\dagger(q)} \int_0^\infty w^\dagger\big(1 - F_P(z)\big)dz$ so that Equation (6) is satisfied by definition. By our choice of q, we have $1 - F_P(z) \leq \frac{q}{w^\dagger(q)} w^\dagger\big(1 - F_P(z)\big)$ for all $z \in [0, \infty)$, and so

$$\int_0^\infty \big(1 - F_P(z)\big)dz \leq \frac{q}{w^\dagger(q)} \int_0^\infty w^\dagger\big(1 - F_P(z)\big)dz = pq. \qquad \square$$

Now we are ready to prove the main result.

Proof of Theorem 4. Let $\mathcal{M}$ be any deterministic-allocation mechanism. We define a deterministic mechanism $\mathcal{M}_\mathcal{I}$ and show that $\text{Rev}_\mathcal{M}(w, \mathcal{D}) \leq \alpha^{-1}\text{Rev}_{\mathcal{M}_\mathcal{I}}(\mathcal{I}, \mathcal{D})$. Thus, $\text{Rev}(w, \mathcal{D}, \mathcal{C}_{da}) \leq \alpha^{-1}\text{DRev}(\mathcal{I}, \mathcal{D}) = \alpha^{-1}\text{DRev}(\mathcal{D})$.

For every $S \subseteq [m]$, let P_S be the corresponding payment variable in $\mathcal{M}$. By Lemma 7, we can assume P_S is equal to p_S with probability q_S and 0 otherwise. In $\mathcal{M}_\mathcal{I}$, we add a corresponding menu item which allocates S with probability 1 and always charges payment $p'_S = p_S(1 - w(1 - q_S))$. Every type v has utility $v(S) - p'_S$ for the menu item in $\mathcal{M}_\mathcal{I}$ that corresponds to S, which is equal to $\mathbb{E}_w[v(S) - P_S]$, the expected utility of type v for the menu item in $\mathcal{M}$ for the same set. Thus, a risk-neutral buyer will purchase in $\mathcal{M}_\mathcal{I}$ the menu item corresponding to what a buyer with weight function w will purchase in $\mathcal{M}$.

It remains to show that p'_S is not too much smaller than $\mathbb{E}[P] = p_S q_S$. By definition of $(\alpha, 0)$-limited, $w(1 - q_S) \leq \alpha(1 - q_S) + 1 - \alpha$. Therefore, $p'_S = p_S(1 - w(1 - q_S)) \geq \alpha q_S p_S$. $\square$

7 Binary Lottery Mechanisms

In this section we study binary lottery mechanisms, the class $\mathcal{C}_b$, as defined in Section 4. Recall that mechanisms in this class are defined such that each menu item is a lottery which allocates some $S \subseteq [m]$ with probability q_S, and charges a payment p_S if and only if S is allocated. This class generalizes the binary lottery mechanisms defined by Chawla et al. [11] for the single item case. Chawla et al. [11] showed that, as long as the buyer is sufficiently pessimistic, the seller can extract nearly the entire expected value of the buyer as revenue, regardless of the buyer's value distribution.

Lemma 8. *[Chawla et al. [11]] For every $\varepsilon > 0$ and $H > 1$, if the buyer's weighting function w is invertible and satisfies $w(1-\varepsilon) \le 2^{-H/\varepsilon}$, there exists a mechanism that for any value distribution $\mathcal{D}$ supported on $[1, H]$ obtains revenue at least $1 - O(\varepsilon)$ times the buyer's expected value $\mathbb{E}[\mathcal{D}]$.*

Observe that under the same conditions we can extract the buyer's full welfare by running the mechanism of Lemma 8 for the grand bundle (which is a valid mechanism for the class $\mathcal{C}_b$). However, as we show next, the revenue is bounded for limited weighting functions.

Theorem 5. *For a single, additive bidder, any $(0, \beta)$-limited weighting function w and any independent distribution $\mathcal{D}$, it holds that*

$$\textsc{Rev}\,(w, \mathcal{D}, \mathcal{C}_b) \le \beta^{-1}\textsc{Rev}(\mathcal{D}).$$

Proof. Let $\mathcal{M} \in \mathcal{C}_b$ be an optimal mechanism for the class $\mathcal{C}_b$. We construct a mechanism $\tilde{\mathcal{M}}$ for a risk-neutral buyer in the following way. Fix v. Let $X(v)$ and $P(v)$ be the allocation and payment that a buyer with type v receives in $\mathcal{M}$. Since this is a binary lottery, let $X(v)$ be supported on $S_v \subseteq [m]$, let $p_v \in \mathbb{R}_{\ge 0}$ be the payment, and let q_v be the probability of allocation. Then the weighted expected utility for this lottery of a buyer with weighting function w and *any* value v' is $u_w(v', X(v), P(v)) = (v'(S_v) - p_v)w(q_v)$. Now, define the lottery $(\tilde{X}(v), \tilde{P}(v))$ as

$$(\tilde{X}(v), \tilde{P}(v)) = \begin{cases} (S_v, p_v) & \text{w.p. } w(q_v) \\ (\emptyset, 0) & \text{o.w.} \end{cases}$$

Then the utility of a risk-neutral buyer with value v' for the lottery $\tilde{X}(v), \tilde{P}(v)$ is also $(v'(S_v) - p_v)w(q_v)$, so that $\tilde{\mathcal{M}}$ is IC and IR for a risk-neutral buyer. The respective revenues are

$$\textsc{Rev}_{\mathcal{M}}(w, \mathcal{D}) = \int_V f_{\mathcal{D}}(v) p_v q_v \, dv \qquad \text{and} \qquad \textsc{Rev}_{\tilde{\mathcal{M}}}(\mathcal{I}, \mathcal{D}) = \int_V f_{\mathcal{D}}(v) p_v w(q_v) \, dv,$$

and so, using the definition of $(0, \beta)$-limited,

$$\frac{\textsc{Rev}_{\tilde{\mathcal{M}}}(\mathcal{I}, \mathcal{D})}{\textsc{Rev}_{\mathcal{M}}(w, \mathcal{D})} \ge \inf \frac{w(q)}{q} \ge \beta.$$

Rearranging, we have $\textsc{Rev}_{\mathcal{M}}(w, \mathcal{D}) \le \beta^{-1}\textsc{Rev}_{\tilde{\mathcal{M}}}(\mathcal{I}, \mathcal{D}) \le \beta^{-1}\textsc{Rev}(\mathcal{I}, \mathcal{D})$. □

Again, combining with Theorem 1 we get the following corollary.

Corollary 4. *For a single, additive bidder, any $(0, \beta)$-limited weighting function w and any independent distribution $\mathcal{D}$, it holds that*

$$\textsc{Rev}\,(w, \mathcal{D}, \mathcal{C}_b) \le 6\beta^{-1} \max\{\textsc{SRev}, \textsc{BRev}\}.$$

References

[1] Maurice Allais. *Fondements d'une théorie positive des choix comportant un risque et critique des postulats et axiomes de l'école américaine*, volume 144. Imprimerie nationale, 1955.

[2] Eduardo M. Azevedo and Daniel Gottlieb. Risk-neutral firms can extract unbounded profits from consumers with prospect theory preferences. *Journal of Economic Theory*, 147(3): 1291 – 1299, 2012. ISSN 0022-0531. doi: https://doi.org/10.1016/j.jet.2012.01.002. URL `http://www.sciencedirect.com/science/article/pii/S0022053112000038`.

[3] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S Matthew Weinberg. A simple and approximately optimal mechanism for an additive buyer. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 21–30. IEEE, 2014.

[4] Graham R Brightwell and Prasad Tetali. The Number of Linear Extensions of the Boolean Lattice. *Order*, 20(4):333–345, 2003.

[5] Yang Cai and Mingfei Zhao. Simple mechanisms for subadditive buyers via duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 170–183. ACM, 2017.

[6] Yang Cai and Mingfei Zhao. Simple mechanisms for subadditive buyers via duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 170–183. ACM, 2017.

[7] Yang Cai, Nikhil R Devanur, and S Matthew Weinberg. A duality-based unified approach to bayesian mechanism design. *ACM SIGecom Exchanges*, 15(1):71–77, 2016.

[8] Colin Camerer. Bounded rationality in individual decision making. *Experimental Economics*, 1(2):163–183, Sep 1998. ISSN 1573-6938. doi: 10.1023/A:1009944326196. URL `https://doi.org/10.1023/A:1009944326196`.

[9] Gabriel Carroll. Robustness and separation in multidimensional screening. *Econometrica*, 85 (2):453–488, 2017.

[10] Shuchi Chawla and J Benjamin Miller. Mechanism design for subadditive agents via an ex ante relaxation. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 579–596. ACM, 2016.

[11] Shuchi Chawla, Kira Goldner, J Benjamin Miller, and Emmanouil Pountourakis. Revenue maximization with an uncertainty-averse buyer. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2050–2068. SIAM, 2018.

[12] Nikhil R Devanur and S Matthew Weinberg. The optimal mechanism for selling to a budget constrained buyer: The general case. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 39–40. ACM, 2017.

[13] Shaddin Dughmi and Yuval Peres. Mechanisms for risk averse agents, without loss. *arXiv preprint arXiv:1206.2957*, 2012.

[14] David Easley and Arpita Ghosh. Behavioral mechanism design: Optimal crowdsourcing contracts and prospect theory. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 679–696. ACM, 2015.

[15] Sebastian Ebert and Philipp Strack. Until the bitter end: On prospect theory in a dynamic context. *American Economic Review*, 105(4):1618–33, April 2015. doi: 10.1257/aer.20130896. URL `http://www.aeaweb.org/articles?id=10.1257/aer.20130896`.

[16] Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, and S Matthew Weinberg. The competition complexity of auctions: A bulow-klemperer result for multi-dimensional bidders. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 343–343. ACM, 2017.

[17] Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, and S Matthew Weinberg. A simple and approximately optimal mechanism for a buyer with complements. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 323–323. ACM, 2017.

[18] Amos Fiat and Christos Papadimitriou. When the players are not expectation maximizers. In *Proceedings of the Third International Conference on Algorithmic Game Theory*, SAGT'10, pages 1–14, Berlin, Heidelberg, 2010. Springer-Verlag. ISBN 3-642-16169-3, 978-3-642-16169-8. URL `http://dl.acm.org/citation.cfm?id=1929237.1929238`.

[19] Hu Fu, Jason Hartline, and Darrell Hoy. Prior-independent auctions for risk-averse agents. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 471–488. ACM, 2013.

[20] Hu Fu, Christopher Liaw, Pinyan Lu, and Zhihao Gavin Tang. The value of information concealment. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2533–2544. Society for Industrial and Applied Mathematics, 2018.

[21] Nick Gravin and Pinyan Lu. Separation in correlation-robust monopolist problem with budget. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2069–2080. SIAM, 2018.

[22] Sergiu Hart and Noam Nisan. Approximate revenue maximization with multiple items. *Journal of Economic Theory*, 172:313–347, 2017.

[23] Daniel Kahneman. Prospect theory: An analysis of decisions under risk. *Econometrica*, 47: 278, 1979.

[24] Xinye Li and Andrew Chi-Chih Yao. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences*, 110(28):11232–11237, 2013.

[25] Siqi Liu and Christos-Alexandros Psomas. On the competition complexity of dynamic mechanism design. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2008–2025. Society for Industrial and Applied Mathematics, 2018. URL `http://dl.acm.org/citation.cfm?id=3174304.3175436`.

[26] Mark J Machina. Choice under uncertainty: Problems solved and unsolved. *Journal of Economic Perspectives*, 1(1):121–154, 1987.

[27] Evdokia Nikolova, Emmanouil Pountourakis, and Ger Yang. Optimal mechanism design with risk-loving agents. In *International Conference on Web and Internet Economics*, pages 375–392. Springer, 2018.

[28] John Quiggin. A theory of anticipated utility. *Journal of Economic Behavior & Organization*, 3(4):323 – 343, 1982. ISSN 0167-2681. doi: https://doi.org/10.1016/0167-2681(82)90008-7. URL `http://www.sciencedirect.com/science/article/pii/0167268182900087`.

[29] John Quiggin. *Generalized expected utility theory: The rank-dependent model*. Springer Science & Business Media, 2012.

[30] Marzena Rostek. Quantile maximization in decision theory. *The Review of Economic Studies*, 77(1):339–371, 2010.

[31] Aviad Rubinstein and S Matthew Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 377–394. ACM, 2015.

[32] David Schmeidler. Subjective probability and expected utility without additivity. *Econometrica: Journal of the Econometric Society*, pages 571–587, 1989.

[33] Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323, 1992. ISSN 1573-0476. doi: 10.1007/BF00122574. URL `http://dx.doi.org/10.1007/BF00122574`.

[34] Peter P. Wakker. *Prospect Theory: For Risk and Ambiguity*. Cambridge University Press, 2010. doi: 10.1017/CBO9780511779329.

[35] John A Weymark. Generalized gini inequality indices. *Mathematical Social Sciences*, 1(4): 409–430, 1981.

[36] Menahem E Yaari. The dual theory of choice under risk. *Econometrica: Journal of the Econometric Society*, 1987.

[37] Andrew Chi-Chih Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pages 92–109. Society for Industrial and Applied Mathematics, 2015.

A Full CPT example

Example 5 (Tversky and Kahneman [33])**.** *Consider the following game of chance. You roll a die once and observe the result* $v = 1, \ldots, 6$*. If* v *is even, you receive* v*; if* v *is odd, you pay* v*. This defines a random variable* X *which takes values* $(-5, -3, -1, 2, 4, 6)$*, each with probability* $1/6$*. Let* X^+ *be the random variable which takes value* 0 *with probability* $1/2$*, and* $2, 4, 6$*, each with probability* $1/6$*. Also, let* X^- *be the random variable which takes the value* 0 *with probability* $1/2$*, and values* $-1, -3, -5$*, each with probability* $1/6$*, correlated such that* $X^+ + X^- = X$*. Assuming that* $r = 0$*, the weighted expectation of* X^+ *is*

$$\mathbb{E}_{w^+}\left[X^+\right] = 2 \cdot (w^+(1/2) - w^+(1/3)) + 4 \cdot (w^+(1/3) - w^+(1/6)) + 6 \cdot (w^+(1/6) - w^+(0)).$$

The intuition here is that the multiplier of value v *is equal to the difference between the weighted probabilities of the events "the outcome of the experiment is at least as good as* v*" and "the outcome is strictly better than* v*". Similarly, the weighted expectation of* X^- *is*

$$\mathbb{E}_{w^-}\left[X^-\right] = (-1) \cdot (w^-(1/2) - w^-(1/3)) + (-3) \cdot (w^-(1/3) - w^-(1/6)) + (-5) \cdot (w^-(1/6) - w^-(0)),$$

where this time the multiplier of v *is equal to the difference between the weighted probabilities of the events "the outcome is at least as bad as* v*" and "the outcome is strictly worse than* v*". Finally, the weighted expectation of* X *is simply* $\mathbb{E}_w[X] = \mathbb{E}_{w^+}[X^+] + \mathbb{E}_{w^-}[X^-]$*.*

B Optimal Max-Min Mechanism

Theorem 6.

$$\max_{\mathcal{M}} \min_{w} \textsc{Rev}_{\mathcal{M}}(w, \mathcal{D}) = \textsc{DRev}(\mathcal{D})$$

Proof. Define w_o, the weighting function of a perfectly optimistic buyer, as the following

$$w_o(q) = \begin{cases} 0 & q = 0 \\ 1 & 0 < q \leq 1 \end{cases}.$$

We first prove the following lemma.

Lemma 9. *For every value distribution* $\mathcal{D}$*, if the buyer's weighting function is* w_o*, then there exists a deterministic mechanism that maximizes the seller's revenue.*

Proof. Consider an arbitrary mechanism $\mathcal{M}$ and a buyer with type $v = (v_1, \ldots, v_m)$. Let $\mathcal{L}$ be a menu item in $\mathcal{M}$. $\mathcal{L}$ defines a distribution over k outcomes: outcome o_i occurs with some probability q_i, where some subset of items S_i is allocated for some payment p_i. Without loss of generality we assume that outcomes are ordered in increasing utility for the buyer. Then the expected utility of the buyer with type v picking $\mathcal{L}$ is

$$\begin{aligned}
\mathbb{E}_{w_o}[v, \mathcal{L}] &= \sum_{i=1}^{k-1} \left(\sum_{j \in S_i} v_j - p_j \right) \left(w_o \left(\sum_{j=i}^{k} q_j \right) - w_o \left(\sum_{j=i+1}^{k} q_j \right) \right) + \left(\sum_{j \in S_k} v_j - p_k \right) w_o(q_k) \\
&= \sum_{j \in S_k} v_j - p_k,
\end{aligned}$$

i.e. the expect utility is the same as the highest utility of all the outcomes in $\mathcal{L}$.

Let $S_{\mathcal{L},v}$ and $p_{\mathcal{L},v}$ be the subset of items and price in the favorite outcome of type v in the menu item $\mathcal{L}$ in $\mathcal{M}$. Now we construct a deterministic mechanism $\mathcal{M}'$: for each type v and menu item $\mathcal{L}$ in $\mathcal{M}$, add to $\mathcal{M}'$ the menu item $\mathcal{L}'_{\mathcal{L},v}$ that deterministically sells $S_{\mathcal{L},v}$ at a price $p_{\mathcal{L},v}$.

It's not hard to see that a buyer with type v and weighting function w_o will buy menu item $\mathcal{L}'_{\mathcal{L},v}$ in $\mathcal{M}'$ if she buys $\mathcal{L}$ in $\mathcal{M}$. Therefore $\textsc{Rev}_{\mathcal{M}}(w_o, \mathcal{D}) = \textsc{Rev}_{\mathcal{M}'}(w_o, \mathcal{D})$. □

Let $\mathcal{M}_{det}$ be the mechanism of Lemma 9. Subsequently, we get that

$$\begin{aligned}
\max_{\mathcal{M}} \min_{w} \textsc{Rev}_{\mathcal{M}}(w, \mathcal{D}) &\leq \max_{\mathcal{M}} \textsc{Rev}_{\mathcal{M}}(w_o, \mathcal{D}) \\
&= \textsc{Rev}_{\mathcal{M}_{det}}(w_o, \mathcal{D}) \\
&\leq \max_{\mathcal{M} \text{ deterministic}} \textsc{Rev}_{\mathcal{M}}(w_o, \mathcal{D}) \\
&= \max_{\mathcal{M} \text{ deterministic}} \textsc{Rev}_{\mathcal{M}}(\mathcal{I}, \mathcal{D}) \\
&= \textsc{DRev}(\mathcal{D})
\end{aligned}$$

The fourth equality holds since a prospect theory buyer has the same preferences as a risk-neutral buyer in a deterministic mechanism. On the other hand:

$$\textsc{DRev}(\mathcal{D}) = \max_{\mathcal{M} \text{ deterministic}} \textsc{Rev}_{\mathcal{M}}(\mathcal{I}, \mathcal{D}) = \max_{\mathcal{M} \text{ deterministic}} \min_{w} \textsc{Rev}_{\mathcal{M}}(w, \mathcal{D}) \leq \max_{\mathcal{M}} \min_{w} \textsc{Rev}_{\mathcal{M}}(w, \mathcal{D}),$$

where the second equality holds since a prospect theory buyer has the same preferences as a risk-neutral buyer in a deterministic mechanism. The theorem follows. □

C Properties of Weighted Expectations

Proof of Lemma 1. We prove the statements for a discrete random variable Z over k outcomes; the proof for continuous random variables is analogous. The i-th outcome in Z, Z_i, occurs with probability p_i, and without loss of generality $Z_i \leq Z_{i+1}$. Notice that for the random variable $W = Z + c$, the ordering remains the same.

$$\begin{aligned}
\mathbb{E}_w[Z + c] &= \sum_{i=1}^{k-1} (Z_i + c) \cdot \left(w\left(\sum_{j=i}^{k} p_j\right) - w\left(\sum_{j=i+1}^{k} p_j\right) \right) + (Z_k + c) w(p_k) \\
&= \sum_{i=1}^{k-1} Z_i \cdot \left(w\left(\sum_{j=i}^{k} p_j\right) - w\left(\sum_{j=i+1}^{k} p_j\right) \right) + Z_k w(p_k) \\
&\quad + \sum_{i=1}^{k-1} c \cdot \left(w\left(\sum_{j=i}^{k} p_j\right) - w\left(\sum_{j=i+1}^{k} p_j\right) \right) + c w(p_k) \\
&= \mathbb{E}_w[Z] + c \cdot \left(\sum_{i=1}^{k-1} w\left(\sum_{j=i}^{k} p_j\right) - w\left(\sum_{j=i+1}^{k} p_j\right) + w(p_k) \right) \\
&= \mathbb{E}_w[Z] + c \cdot w\left(\sum_{j=1}^{k} p_j\right)
\end{aligned}$$

$$= \mathbb{E}_w[Z] + c.$$

Similarly,

$$\begin{aligned}
\mathbb{E}_w[cZ] &= \sum_{i=1}^{k-1}(cZ_i) \cdot \left(w\left(\sum_{j=i}^{k} p_j\right) - w\left(\sum_{j=i+1}^{k} p_j\right)\right) + (cZ_k)w(p_k) \\
&= c \cdot \left(\sum_{i=1}^{k-1} Z_i \cdot \left(w\left(\sum_{j=i}^{k} p_j\right) - w\left(\sum_{j=i+1}^{k} p_j\right)\right) + Z_k w(p_k)\right) \\
&= c \cdot \mathbb{E}_w[Z] \,.
\end{aligned}$$

□

Proof of Claim 2.

$$\begin{aligned}
\mathbb{E}_w[-X] &= -\int_{-\infty}^{0} \big(1 - w(1 - F_{-X}(z))\big)dz + \int_{0}^{\infty} w(1 - F_{-X}(z))dz \\
&= -\int_{0}^{\infty} \big(1 - w(F_X(z))\big)dz + \int_{-\infty}^{0} w(F_X(z))dz \\
&= -\mathbb{E}_{w^\dagger}[X]
\end{aligned}$$

□