# Playing Divide-and-Choose Given Uncertain Preferences 

Jamie Tucker-Foltz ${ }^{1}$ and Richard Zeckhauser ${ }^{2}$<br>${ }^{1}$ Harvard University, jtuckerfoltz@gmail.com<br>${ }^{2}$ Harvard University, richard_zeckhauser@harvard.edu

January 1, 2024


#### Abstract

We study the classic divide-and-choose method for equitably allocating divisible goods between two players who are rational, self-interested Bayesian agents. The players have additive values for the goods. The prior distributions on those values are common knowledge. We consider both the cases of independent values and values that are correlated across players (as occurs when there is a common-value component).

We describe the structure of optimal divisions in the divide-and-choose game and identify several cases where it is possible to efficiently compute equilibria. An approximation algorithm is presented for the case when the distribution over the chooser's value for each good follows a normal distribution, along with a randomized approximation algorithm for the case of uniform distributions over intervals.

A mixture of analytic results and computational simulations illuminates several striking differences between optimal strategies in the cases of known versus unknown preferences. Most notably, given unknown preferences, the divider has a compelling "diversification" incentive in creating the chooser's two options. This incentive leads to multiple goods being divided at equilibrium, quite contrary to the divider's optimal strategy when preferences are known.

In many contexts, such as buy-and-sell provisions between partners, or in judging fairness, it is important to assess the relative expected utilities of the divider and chooser. Those utilities, we show, depend on the players' levels of knowledge about each other's values, the correlations between the players' values, and the number of goods being divided. Under fairly mild assumptions, we show that the chooser is strictly better off for a small number of goods, while the divider is strictly better off for a large number of goods.


## 1 Introduction

Ever since Abraham and Lot divided the land of Canaan, with Abraham dividing and Lot choosing, the divide-and-choose method has been employed to parcel out assets. Today the method is widely used when partners in a real estate deal invoke their buy-sell agreement, or when siblings divide up an inheritance.

The assets can be continuous, as with the classic cake-cutting problem, or separable and indivisible, for example valuables from an estate. Consider multi-dimensional cake-cutting with divisible elements vanilla filling, chocolate icing and a cherry. One player will divide the cake; the other will then choose her piece. Positing additive preferences and that fractional portions have fractional values, if the divider knows the chooser's preferences, he arrays the assets (filling, icing, cherry) in order of the ratio of his value to her (the chooser's) value. Just enough low-ratio assets are placed
in pile 2 to assure that the chooser selects it, leaving the divider the high-ratio assets in pile $1 .{ }^{1}$ Only one asset will generally need to be divided fractionally; in knife-edge situations it could be none.

Now consider a more consequential example: Partners in a real-estate company have signed a buy-sell arrangement. As is common, after a specified period of years, either partner can trigger the arrangement by creating a two-pile division. Pile 1 might consist of asset $Q$ and a required payment of $\$ 1.5$ million. Pile 2 would contain the remaining asset, $R$, and receipt of $\$ 1.5$ million. Partner A launches as the divider. Partner B must then choose between the two piles. The very nature of the legal agreement allows A to reject any proposals from B to adjust terms.

In real life, the players rarely know each others' preferences. The cake divider will know his own preferences, but will only have a feel-that is, a Bayesian prior-for how the chooser values filling to icing to cherry, or asset $Q$ to asset $R$ to money. While the divider can think in terms of expected ratios, he will remain uncertain of the chooser's total value for any two piles. The divider's optimal strategy will weigh the disadvantages of putting fewer goods (or lesser amounts of divisible goods) in pile 1 against the likelihood that the chooser picks pile 1. How that trade-off should be handled is the subject of this paper.

### 1.1 Our contributions

Our central contribution is to analyze the divide-and-choose game when information is asymmetric, i.e., when the players' values are private information, but priors for those values are common knowledge. This is the typical setup for self-interested players in Bayesian games. Section 2 presents our model for divide-and-choose.

The literature on divide-and-choose games focuses overwhelmingly on the case of known preferences, though real-world situations rarely meet that standard. In the comforting land of known preferences, the divider follows the simple optimization procedure mentioned above. In sharp contrast, once uncertainty enters, a strategic Bayesian divider generally faces a complicated optimization problem, potentially with a myriad of local optima. In Section 3 we show that such complexities can arise even for extremely simple priors, such as independently and identically distributed (i.i.d.) normal distributions.

The apparent lack of structure to the divider's optimization problem suggests that it may be computationally intractable in general. One of our main contributions is an algorithmic framework for optimizing over divisions that yields polynomial-time approximation algorithms when the distributions over the chooser's values for each good satisfy certain properties. Specifically, when each distribution is symmetric about its mean and log-concave, we show that there exists a sequence of convex programs computing divisions that yield expected utility arbitrarily close to the maximum possible expected utility (Theorem 3.9). As special cases, we obtain a fully polynomial-time approximation scheme (FPTAS) when the distributions are Gaussian (Theorem 3.6) and a fully polynomial-time randomized approximation scheme (FPRAS) when the distributions are uniform over real intervals (Theorem 3.10). We also present an algorithm that solves the case of discrete distributions, which is practical when the number of possible chooser types is small.

This complexity of the divider's optimization problem arises because, in contrast to the case of known preferences, it is often beneficial to divide more than one good. To take our cake analogy, if a divider prefers filling and suspects that the chooser has a high value for both the icing and the cherry, pile 1 may optimally consist of $100 \%$ of the filling, plus $59 \%$ of the icing and $31 \%$ of the cherry. By ensuring that pile 2 contains a bit of icing and cherry, he decreases the variance

[^0]in the chooser's valuation for pile 1, hence reducing the likelihood that she will choose that pile. In Section 4 we investigate the divider's incentive to diversify his risk by dividing several goods (or, for indivisible assets, committing to lotteries that randomly allocate some goods between the piles after the chooser picks a pile). Even for a risk-neutral divider, diversification is warranted for a wide range of prior distributions. This is still the case even when the divider's values are strongly correlated with those of the chooser. Furthermore, we show that a risk averse divider should diversify more. These results still do not determine which goods should be split between the two piles. Simple rules, we show, can lead the divider astray.

Finally, in Section 5 we analyze the expected utilities of the divider and chooser under a range of circumstances. For example, which player is better off a priori if both players' values are drawn from the same distribution? With known preferences, it is far preferable to be the divider. With significant uncertainties about the players' values of the individual goods, the chooser is better off when there are few goods, but the advantage tips to the divider as the number of goods increases.

Our Conclusion, Section 6, presents a number of open questions that emerged from this analysis. It closes by stressing that although our focus has been on theory and computational methods, the analysis is widely applicable in real-world contexts. The divide-and-choose method, or close analogues, though rarely identified that way, are widely used in practice. Take-it-or-leave-it offers and buy-and-sell agreements between business partners are salient examples.

### 1.2 Related work

The divide-and-choose method features prominently in the literature on fair division. In the cakecutting model [25], agents are assumed to have additive, divisible preferences over the unit interval $[0,1]$ (the "cake") and a feasible allocation is a partitioning of $[0,1]$ between the players, where each player's part of the allocation is a finite union of intervals. With two players, the divide-and-choose method is particularly useful in this model for the following reasons:
(i) If the first player divides the cake equally, the allocation will be envy-free, meaning that the players each value their own piece of cake higher.
(ii) The protocol can be implemented by asking the players simple queries.

Extensions have also been discovered for 3 or more players [8, 2]. These methods require several rounds of dividing and choosing, but ultimately satisfy the same two properties. Kuhn [19] gives abstract axioms under which a simple one-round variant of divide-and-choose with one divider and $n-1$ choosers will yield an acceptable outcome for all $n$ players. Another major line of work studies deterministic allocations of indivisible goods, in which case envy-freeness is sometimes not feasible. Think of three goods each valued roughly the same by both players. They can at best be divided 1 and 2 , so the divider will generally be envious. Hence, various natural relaxations have been studied $[10,14]$. Variant settings with more complex kinds of goods and constraints have also been studied, such as goods with unknown, random values [6].

The fair division literature has largely focused on finding mechanisms satisfying axiomatic properties such as envy-freeness, efficiency, etc., when players do not behave strategically. A handful of works consider the objective of strategy-proofness as well, both for divisible [13, 12, 4, 5, 11] and indivisible $[3,1]$ goods. These works largely draw on the mechanism design perspective, studying questions such as, what is the optimal worst-case utility approximation to the socially optimal outcome under the constraint of strategy-proofness?

Our work is more in the spirit of Clausewitz the strategist than Solomon the arbitrator: Our prime interest is not in finding mechanisms to implement socially desirable outcomes. Rather, we
focus on how strategic players should actually behave to maximize their personal welfare. Issues of welfare and efficiency enter our analysis, particularly how those issues are affected by some simple extensions of the divide-and-choose game. However, our goal is not to modify the game to bolster its fairness or efficiency, but rather to analyze the effects from a descriptive perspective. There is a small experimental literature on strategic behavior under various cake-cutting protocols; for instance, see Kyropoulou, Ortega, and Segal-Halevi [20]. On the theoretical front, however, much less is known. We are aware of one prior work, by Delgosha and Gohari [15], that studies optimal strategies in an environment where players may learn about each other's preferences through repeated interactions. Our interest is in the more common setting of one-shot interactions.

A key concept we discuss is the critical ratio of a good, which we define to be the ratio of the divider's value to the expectation of the chooser's value. This is a ubiquitous notion in contexts from hypothesis testing to cost/effectiveness analyses. (See, for example, Weinstein and Zeckhauser [26].) In the case where the divider has perfect knowledge of the chooser's preferences, the divider places the goods with the highest critical ratio in pile 1 and those with the lowest critical ratio in pile 2 . He determines the cutoff so that the chooser just picks pile 2. The optimality of this strategy in the setting of cake-cutting is noted by H. P. Young [27], as well as Brânzei, Caragiannis, Kurokawa, and Procaccia [9]. One of the key conceptual takeaways from this paper is that uncertainty over preferences notably complicates the task of computing optimal divisions in a way that is not characterized by critical ratios.

Finally, we note that our model bears a technical resemblance to a model of final-offer arbitration (FOA) by Powers [24] and Brams and Merrill [7]. Just as in our setting, there are multiple divisible goods that must be allocated among two players. However, instead of playing divide-and-choose, both players simultaneously submit divisions to an arbiter, who chooses the more "reasonable" one, where players have common knowledge of the arbiter's sense of reasonableness. A key technical difference between FOA and divide-and-choose is that, in FOA, the uncertainty is not over the other player's values, but over the other player's action. Yet both models involve similar optimal strategies: give just enough away to the other player to obtain the desired binary decision with decent probability.

## 2 Model

There are two players: the divider $(D)$ and the chooser $(C)$. There are $n$ divisible goods, which we number from 1 to $n$, writing $[n]:=\{1,2,3, \ldots, n\}$ for the set of goods. We may also think of the goods as indivisible, in which case we allow the divider to fractionally allocate goods between the piles via lotteries that are resolved after the chooser has chosen her pile. A player's value is simply the sum of the values of the goods received. Players are risk-neutral, except in Section 4.3 on risk aversion. For each $i \in[n]$, we denote the respective private values of good $i$ to the divider and chooser by $g_{i}^{D}$ and $g_{i}^{C}$, which are drawn from a joint distribution $\mathcal{G}_{i}$ over $\mathbb{R}^{2}$. We write $\mathcal{G}_{i}^{D}$ and $\mathcal{G}_{i}^{C}$ for the respective marginal distributions. More compactly, we denote the value vectors $g^{D}:=\left(g_{1}^{D}, g_{2}^{D}, \ldots, g_{n}^{D}\right)$ and $g^{C}:=\left(g_{1}^{C}, g_{2}^{C}, \ldots, g_{n}^{C}\right)$, which are drawn according to the joint distribution $\mathcal{G}=\mathcal{G}_{1} \times \mathcal{G}_{2} \times \cdots \times \mathcal{G}_{n}$. Thus, we allow for values of the two players for any individual good to be correlated, but require independence across goods. We write the respective marginal distributions $\mathcal{G}^{D}:=\mathcal{G}_{1}^{D} \times \mathcal{G}_{2}^{D} \times \cdots \times \mathcal{G}_{n}^{D}$ and $\mathcal{G}^{D}:=\mathcal{G}_{1}^{C} \times \mathcal{G}_{2}^{C} \times \cdots \times \mathcal{G}_{n}^{C}$. We write $\overline{\mathcal{G}}_{i}^{C}$ and $\overline{\mathcal{G}}^{C}$ for the updated distributions over the chooser's values after the divider observes his values. The divide-and-choose game proceeds in three steps.
(i) Values are drawn $\left(g^{D}, g^{C}\right) \sim \mathcal{G}$. The divider observes $g^{D}$ and the chooser observes $g^{C}$.
(ii) The divider chooses a division of the goods, which is a vector $p=\left(p_{1}, p_{2}, \ldots, p_{n}\right) \in[0,1]^{n}$. We refer to pile 1 as the allocation consisting of $p_{i}$ of each good $i$ and pile 2 as the allocation consisting of $\left(1-p_{i}\right)$ of each good $i$.
(iii) The chooser picks her higher-valued pile, namely pile 1 or pile 2 , for herself. The other pile goes to the divider.
Formally, the payoffs are defined as follows. If the chooser picks pile 1 , then the divider receives a payoff of $u^{D}=\sum_{i=1}^{n}\left(1-p_{i}\right) g_{i}^{D}$ and the chooser receives a payoff of $u^{C}=\sum_{i=1}^{n} p_{i} g_{i}^{C}$. On the other hand, if the chooser picks pile 2 , then the divider receives a payoff of $u^{D}=\sum_{i=1}^{n} p_{i} g_{i}^{D}$ and the chooser receives a payoff of $u^{C}=\sum_{i=1}^{n}\left(1-p_{i}\right) g_{i}^{C}$. To facilitate understanding, we refer to pile 1 as the pile that the divider would prefer to have for himself, an assumption that is without loss of generality (see Lemma 3.2). Given a division $p$, we refer to the probability that the chooser picks pile 1 as $P$.

Thus, the game is completely parameterized by the value distribution $\mathcal{G}$. In this paper we mostly focus on settings where the values are independent across players, except in Sections 4.2 and 5.1. We especially focus on the following distributions:

- Normal priors: The value for good each $i$ is drawn from $\mathcal{N}\left(\mu_{i}, \sigma_{i}\right)$ for some mean $\mu_{i} \in \mathbb{R}$ and standard deviation $\sigma_{i}>0$.
- Discrete priors: The value for each good $i$ is drawn from an arbitrary distribution over $\mathbb{R}$ with finite support.
- Uniform priors: The value for each good $i$ is drawn i.i.d. from uniform distributions over intervals $\left[a_{i}, b_{i}\right.$ ], where $0 \leq a_{i} \leq b_{i}$.
Throughout this paper we often assume that each $g_{i}^{D}$ has been fixed, which makes the distribution $\mathcal{G}_{i}^{D}$ irrelevant. Given fixed values of $g_{i}^{D}$, we define the critical ratio of good $i$ to be

$$
C R_{i}:=\frac{g_{i}^{D}}{\mathbb{E}\left[g_{i}^{C}\right]},
$$

where the expectation is over $\overline{\mathcal{G}}_{i}$, i.e., the conditional expectation of $g_{i}^{C}$ given $g_{i}^{D}$.
Note that our model allows for goods to take potentially negative values (i.e., be "bads"). With normal priors this will happen with some nonzero probability, though in most of our examples the means are sufficiently large to render this probability negligibly small. For some of our results we also require that values be nonnegative. For instance, we only discuss critical ratios in the context where all divider values and chooser means are positive.

Our solution concept is that of a subgame-perfect, Bayes-Nash equilibrium. Without loss of generality we need only consider pure strategies for both players, since at no point can either player benefit from inferring information about the other player's type from its actions. We make a minor additional technical assumption: indifference is always broken in favor of the other player. This is only necessary for knife-edge cases in some of our theorems. In practice, we would expect optimal divisions and optimal pile choices to be unique, rendering this assumption unnecessary.

## 3 Computing optimal divisions

The divider first observes his values for the $n$ goods $g_{1}^{D}, g_{2}^{D}, \ldots, g_{n}^{D}$. How should he use that information together with the updated priors ${ }^{2} \overline{\mathcal{G}}_{1}^{C}, \overline{\mathcal{G}}_{2}^{C}, \ldots, \overline{\mathcal{G}}_{n}^{C}$ on the chooser's values to determine

[^1]his optimal division? We first establish some basic, general properties of the equilibria of the divide-and-choose game. We then describe algorithms to determine the divider's optimal division for both normal priors and discrete priors. Finally, we generalize our algorithm for normal priors to handle a wider class of distributions, one that includes uniform priors.

### 3.1 The general structure of optimal divisions

We begin by recording two easy facts about optimal divisions, which we will use throughout the paper. Their straightforward proofs are deferred to Appendices A and B.

The first observation, to use terminology from the fair division literature, says that all equilibria satisfy proportionality for the chooser and expected proportionality for the divider. In other words, both players can expect to take away at least half of their total values.

Lemma 3.1. In any equilibrium of the divide-and-choose game, the following hold.
(i) For any realizations of $g^{D}$ and $g^{C}, u^{C} \geq \frac{1}{2} \sum_{i=1}^{n} g_{i}^{C}$.
(ii) For any realization of $g^{D}, \mathbb{E}_{g C \sim \overline{\mathcal{G}}^{C}}\left[u^{D}\right] \geq \frac{1}{2} \sum_{i=1}^{n} g_{i}^{D}$.

We refer to the utilities $\frac{1}{2} \sum_{i=1}^{n} g_{i}^{D}$ and $\frac{1}{2} \sum_{i=1}^{n} g_{i}^{C}$ as the respective proportionality guarantees utilities of the divider and chooser.

It is convenient to consider the divider's optimization problem, not employing the $p_{i}$ variables taking values in $[0,1]$, but instead in terms of the auxiliary variables

$$
\begin{equation*}
q_{i}:=2 p_{i}-1=p_{i}-\left(1-p_{i}\right) \tag{1}
\end{equation*}
$$

taking values in $[-1,1]$. For reference, the inverse of this correspondence is

$$
\begin{equation*}
p_{i}=\frac{q_{i}}{2}+\frac{1}{2} \tag{2}
\end{equation*}
$$

We frequently refer to a division as $p$ or $q$ interchangeably.
Lemma 3.2. A divider-optimal division with the following two properties always exists.

- The divider weakly prefers pile 1:

$$
\begin{equation*}
\sum_{i=1}^{n} q_{i} g_{i}^{D} \geq 0 \tag{3}
\end{equation*}
$$

- The chooser is weakly more likely to pick pile 2:

$$
\begin{equation*}
P:=\operatorname{Pr}_{g^{C} \sim \overline{\mathcal{G}}^{C}}\left[\sum_{i=1}^{n} q_{i} g_{i}^{C} \geq 0\right] \leq \frac{1}{2} \tag{4}
\end{equation*}
$$

Furthermore, the divider can achieve a strictly higher interim expected utility than his proportionality guarantee if and only if both of these inequalities can be made strict.

For a wide class of priors over the chooser's value, including normal and uniform distributions, we can characterize exactly when the inequalities from Lemma 3.2 will be strict.

Proposition 3.3. Suppose each $\overline{\mathcal{G}}_{i}^{C}$ is a non-atomic distribution that is symmetric about its mean, and suppose all divider values and chooser means are positive. Then the optimal division yields utility equal to the proportionality guarantee if and only if all goods have the same critical ratios.

Proof. Suppose first that all goods have the same critical ratio $r$. First note that the division $q_{1}=q_{2}=\cdots=q_{n}=0$ For any other division $q$, where at least one $q_{i}$ is not zero, consider the random variable

$$
\begin{equation*}
X:=\sum_{i=1}^{n} q_{i} g_{i}^{C} \tag{5}
\end{equation*}
$$

Since the distribution of $X$ is a nontrivial linear combination of the $\overline{\mathcal{G}}_{i}^{C}$, it is also non-atomic and symmetric about its mean, which implies that, for any threshold $t$,

$$
\begin{equation*}
\operatorname{Pr}[X>t]>\frac{1}{2} \Longleftrightarrow \mathbb{E}[X]>T \tag{6}
\end{equation*}
$$

Therefore,

$$
\begin{aligned}
P>\frac{1}{2} & \Longleftrightarrow \operatorname{Pr}\left[\sum_{i=1}^{n} q_{i} g_{i}^{C}>0\right]>\frac{1}{2}
\end{aligned} \begin{aligned}
& \Longleftrightarrow \mathbb{E}\left[\sum_{i=1}^{n} q_{i} g_{i}^{C}\right]>0 \quad(\text { from }(6)) \\
& \\
& \\
& \\
& \left.\Longleftrightarrow \sum_{i=1}^{n} q_{i} g_{i}^{D} \frac{g_{i}^{C}}{g_{i}^{D}}\right]>0
\end{aligned} \begin{aligned}
& i=1 \\
& \\
&
\end{aligned}
$$

In other words, the chooser is strictly more likely to pick pile 1 if and only if the divider strictly prefers pile 1. That it turn implies that it is not possible to achieve a higher-than-proportional utility by the final statement of Lemma 3.2.

Now suppose that two goods $j$ and $k$ have different critical ratios, i.e., $\overline{\mathcal{G}}_{j}^{C}$ has mean $\mu_{j}$ and $\overline{\mathcal{G}}_{k}^{C}$ has mean $\mu_{k}$ such that,

$$
\frac{g_{j}^{D}}{\mu_{j}}>\frac{g_{k}^{D}}{\mu_{k}}
$$

Let $\alpha:=\max \left\{g_{j}^{D}+\mu_{j}, g_{k}^{D}+\mu_{k}\right\}$, let $q$ be the division such that

$$
\begin{aligned}
q_{j} & :=\frac{g_{k}^{D}+\mu_{k}}{\alpha} \\
q_{k} & :=-\frac{g_{j}^{D}+\mu_{j}}{\alpha} \\
q_{i} & :=0 \quad(\text { for all } i \in[n] \backslash\{j, k\})
\end{aligned}
$$

Note that $q_{i}=0$ corresponds to $p_{i}=\frac{1}{2}$, so in words, this is a division where good $j$ is slightly more in pile 1 , good $k$ is slightly more in pile 2 , and all other goods are divided equally between the two piles (scaling by $\alpha$ ensures $\left|q_{j}\right|,\left|q_{k}\right| \leq 1$ ). Then

$$
\sum_{i=1}^{n} q_{i} g_{i}^{D}=q_{j} g_{j}^{D}+q_{k} g_{k}^{D}=\frac{g_{k}^{D}+\mu_{k}}{\alpha} g_{j}^{D}-\frac{g_{j}^{D}+\mu_{j}}{\alpha} g_{k}^{D}=\frac{\mu_{k} g_{j}^{D}-\mu_{j} g_{k}^{D}}{\alpha}>0
$$

so the divider strictly prefers pile 1 . Also, the random variable $X$ has mean

$$
\sum_{i=1}^{n} q_{i} \mu_{i}=q_{j} \mu_{j}+q_{k} \mu_{k}=\frac{g_{k}^{D}+\mu_{k}}{\alpha} \mu_{j}-\frac{g_{j}^{D}+\mu_{j}}{\alpha} \mu_{k}=\frac{g_{k}^{D} \mu_{j}-g_{j}^{D} \mu_{k}}{\alpha}<0
$$

which implies (via (6)) that

$$
P=\operatorname{Pr}\left[\sum_{i=1}^{n} q_{i} g_{i}^{C}>0\right]<\frac{1}{2} .
$$

Therefore, it follows from the final statement of Lemma 3.2 that the divider achieves a utility that is higher than his proportionality guarantee.

We remark that this result does not extend to some other natural families of distributions, for instance, discrete 2-point distributions. In fact, even when all chooser priors and divider values are identical, there may exist bizarre "symmetry-breaking" divisions that yield utility higher than the divider's proportionality guarantee; see Proposition 4.1.

### 3.2 Normal priors

We now turn to prove the first main result of this paper: an efficient algorithm to compute a near-optimal division given the divider's values for each good and normal priors for the chooser values. The procedure is presented formally as Algorithm 1; first we give an informal explanation. With a bit of manipulation, we can rewrite the divider's optimization problem as maximizing

$$
\mathbb{E}\left[u^{D}\right]=\sum_{i=1}^{n} \frac{g_{i}^{D}}{2}\left(P\left(1-q_{i}\right)+(1-P)\left(1+q_{i}\right)\right)
$$

over the variables $q_{1}, q_{2}, \ldots, q_{n} \in[-1,1]$. This is not a linear program. However, the only nonlinearities arise from the $P$ term, which is itself a function of the $q_{i}$. The key idea is to try to guess the optimal $P$ by trying several different values, uniformly spread out between 0 and $\frac{1}{2}$. For each guessed value of $P$, we add a constraint that the chooser picks pile 1 with probability at most $P$. For normal priors, this turns the linear program into a quadratic program, yet it is still convex, so can be readily solved in polynomial time. The only catch is that we lose exact optimality, picking up an error term from potentially missing the exact optimal value of $P$. Fortunately, we can use the structure of the objective function to bound this error in a way that gives a very strong guarantee of approximate optimality.

In what follows, $\Phi$ denotes the standard normal cumulative distribution function.
Lemma 3.4. Algorithm 1 runs in polynomial time in the values of $n$ and $\frac{\sum_{i=1}^{n}\left|g_{i}^{D}\right|}{\gamma}$.
Proof. Observe that there are at most

$$
\frac{1}{2 \delta}=\frac{\sum_{i=1}^{n}\left|g_{i}^{D}\right|}{2 \gamma}
$$

iterations of the main loop. Thus, all that remains to show is that each iteration takes polynomial time in $n$. This follows from the observation that each $\mathcal{C}_{P}$ is a convex program. To see this, note that the objective function is clearly linear in the $q_{i}$ variables, and all constraints except for the final one are linear as well. The final constraint is not linear, but we claim that the set $S \subseteq \mathbb{R}^{n}$ of points satisfying the constraint is convex. Suppose $q:=\left(q_{1}, q_{2}, \ldots, q_{n}\right) \in S$, and let $\widetilde{q}:=$ $\left(\sigma_{1} q_{1}, \sigma_{2} q_{2}, \ldots \sigma_{n} q_{n}\right)$. Then, for any positive real number $c$, the scaled vector $c q=\left(\frac{q_{1}}{c}, \frac{q_{2}}{c}, \ldots, \frac{q_{n}}{c}\right)$
lies in $S$ as well, since

$$
\begin{aligned}
\sum_{i=1}^{n} \mu_{i}\left(c q_{i}\right)=c \cdot \sum_{i=1}^{n} \mu_{i} q_{i} & \leq c \cdot \Phi^{-1}(P) \sqrt{\sum_{i=1}^{n} \sigma_{i}^{2} q_{i}^{2}} \quad(\text { since } q \in S) \\
& =\Phi^{-1}(P) c\|\widetilde{q}\|_{2} \quad\left(\text { where }\|\cdot\|_{2}\right. \text { denotes the Euclidean norm) } \\
& =\Phi^{-1}(P)\|c \widetilde{q}\|_{2} \\
& =\Phi^{-1}(P) \sqrt{\sum_{i=1}^{n} \sigma_{i}^{2}\left(c q_{i}\right)^{2}}
\end{aligned}
$$

Therefore, $S$ is a cone centered at the origin, which is a convex set.

```
Algorithm 1: Computes an approximately optimal division given the divider's values and
independent normal priors for the chooser's values.
    Input: Divider values \(g_{1}^{D}, g_{2}^{D}, \ldots, g_{n}^{D}\) (not all zero), prior means \(\mu_{1}, \mu_{2}, \ldots, \mu_{n}\),
        corresponding standard deviations \(\sigma_{1}, \sigma_{2}, \ldots, \sigma_{n}\), and an additive error bound \(\gamma\)
        for the divider's optimal utility.
    Output: An approximately optimal division \(p_{1}, p_{2}, \ldots, p_{n}\).
    \(\delta \leftarrow \frac{\gamma}{\sum_{i=1}^{n}\left|g_{i}^{D}\right|}\)
    \(P \leftarrow \frac{1}{2}\)
    \(u \leftarrow-\infty\)
    while \(P>0\) do
    \(u_{P}, q_{1}, q_{2}, \ldots, q_{i} \leftarrow\) optimal solution to the following program \(\mathcal{C}_{P}\), which has variables
        \(q_{1}, q_{2}, \ldots, q_{i}:\)
            maximize \(\quad u_{P}=\sum_{i=1}^{n} \frac{g_{i}^{D}}{2}\left(P\left(1-q_{i}\right)+(1-P)\left(1+q_{i}\right)\right)\)
                    \(-1 \leq q_{i} \leq 1 \quad\) for all \(1 \leq i \leq n\),
                        \(\sum_{i=1}^{n} g_{i}^{D} q_{i} \geq 0\),
                        \(\sum_{i=1}^{n} \mu_{i} q_{i} \leq \Phi^{-1}(P) \sqrt{\sum_{i=1}^{n} \sigma_{i}^{2} q_{i}^{2}}\)
        if \(u_{P}>u\) then
        \(u \leftarrow u_{P}\)
        for \(i \leftarrow 1,2, \ldots, n\) do
            \(p_{i} \leftarrow \frac{q_{i}+1}{2}\)
        end
    end
    \(P \leftarrow P-\delta\)
    end
    return \(\left(p_{1}, p_{2}, \ldots, p_{n}\right)\)
```

We remark that, beyond being efficient in theory, this algorithm proves to be fast in practice. We implemented this algorithm using Gurobi [18] to solve $\mathcal{C}_{P}$ as a convex quadratic program and used it to verify many of the examples in this paper.

Lemma 3.5. Algorithm 1 finds a division yielding divider utility within an additive $\gamma$ of the optimal divider utility.

The proof is long and technical, so it is deferred to Appendix C.
With a bit more work, we can translate this additive approximation guarantee into a multiplicative one. Formally, we can show that Algorithm 1 is a fully polynomial-time approximation scheme (FPTAS) for maximizing divider utility, which means that, on instances where the optimal value is $u^{*}>0$, for any $\varepsilon>0$, it can find a solution with objective value at least $(1-\varepsilon) \cdot u^{*}$ in time polynomial in both $n$ and $\frac{1}{\varepsilon}$.
Theorem 3.6. When all $g_{i}^{D} \geq 0$, running Algorithm 1 with $\gamma:=\frac{\varepsilon}{2} \cdot \sum_{i=1}^{n} g_{i}^{D}$ is a fully polynomialtime approximation scheme with approximation parameter $\varepsilon$.

Proof. Given this choice of $\gamma$, it follows from Lemma 3.4 and the fact that $\left|g_{i}^{D}\right|=g_{i}^{D}$ that the algorithm runs in polynomial time in $n$ and $\frac{1}{\varepsilon}$. Let $u$ denote the utility of the solution returned by the algorithm, and let $u^{*}$ denote the optimal utility. By Lemma 3.1, $2 u^{*} \geq \sum_{i=1}^{n} g_{i}^{D}$. Combining this inequality with Lemma 3.5 , we have

$$
u \geq u^{*}-\gamma=u^{*}-\frac{\varepsilon}{2} \cdot \sum_{i=1}^{n} g_{i}^{D} \geq u^{*}-\frac{\varepsilon}{2} \cdot 2 u^{*}=(1-\varepsilon) \cdot u^{*} .
$$

as desired.
One might wonder why it is necessary to sequentially search for the optimal value of $P$. If the optimal divider utility given $P$ were a single-peaked function of $P$, then it would be possible to rapidly compute an optimal division through a ternary search over $P$. However, single-peakedness is not always satisfied, as Figure 1 demonstrates. In fact, local optima can occur even in very simple scenarios:

Proposition 3.7. Even for $n=3$ goods, there exist positive divider values and normal chooser priors with identical positive means such that the divider's interim expected utility, as a function of the division p, has a local maximum that is not a global maximum.

Proof. This was discovered and verified via computational methods, so here we just explain the example at a high level. Suppose $\overline{\mathcal{G}}_{1}^{C}=\mathcal{N}(100,1), \overline{\mathcal{G}}_{2}^{C}=\mathcal{N}(100,1), \overline{\mathcal{G}}_{3}^{C}=\mathcal{N}(100,4225)$, and $g_{1}^{D}=11, g_{2}^{D}=9, g_{3}^{D}=1$. There are two locally-optimal strategies. Approximately, they are:
(i) Divide the high-variance good 3 evenly between the two piles, so that it has no influence on the probability that the chooser picks pile 1 . Then execute the optimal perfect-information strategy, putting most of good 1 in pile 1 and all of good 2 in pile 2 . The risk that the chooser picks pile 1 is very low, at $P \approx 0.015$.
(ii) Put good 3 entirely into pile 2 , so that it is possible to extract a substantial amount of both goods 1 and 2 in pile 1 . The risk that the chooser picks pile 1 is moderate, at $P \approx 0.21$.

As we verified using Algorithm 1, strategy (i) yields utility of approximately 11 (as one would expect), whereas, despite the risk in relying on the high-variance good, strategy (ii) yields utility of approximately 12. However, intermediate strategies yield utilities less than either one of these extremes.


Figure 1: An instance with four goods where there are four locally-optimal divisions with a variety of different values of $P$. The globally optimal value of $P$ is indicated as $P^{*}$. The divider values are $g_{1}^{D}=3, g_{2}^{D}=2, g_{3}^{D}=1, g_{4}^{D}=1.2$, and corresponding chooser priors are $\overline{\mathcal{G}}_{1}^{C}=\mathcal{N}(5,1)$, $\overline{\mathcal{G}}_{2}^{C}=\mathcal{N}(9.5,1), \overline{\mathcal{G}}_{3}^{C}=\mathcal{N}(13.6,96.04), \overline{\mathcal{G}}_{4}^{C}=\mathcal{N}(95,28561)$.

### 3.3 Symmetric log-concave priors

It turns out that the technique from Algorithm 1 applies to a much wider class of distributions. Specifically, we consider distributions $\mathcal{D}$ whose density functions have two properties: symmetry about their mean and log-concavity. While we are not able to get a clean quadratic program, as we obtain in the case of normal distributions, we can still show convexity. As an application, we give a randomized approximation algorithm that handles the case where each $\overline{\mathcal{G}}_{i}^{C}$ is a uniform distribution over an interval.

Many well-studied families of continuous distributions are symmetric and log-concave, including normal distributions and uniform distributions over intervals. A product measure

$$
\mathcal{D}=\mathcal{D}_{1} \times \mathcal{D}_{2} \times \cdots \times \mathcal{D}_{n}
$$

is symmetric and $\log$-concave in an $n$-dimensional sense whenever each $\mathcal{D}_{i}$ is. Our interest in these distributions stems from the following theorem.

Theorem 3.8 (Meyer and Reisner [22]). Let $\mathcal{D}$ be a symmetric, log-concave measure over $\mathbb{R}^{n}$ and let $0 \leq P \leq \frac{1}{2}$. There exists a set $K_{P} \subseteq \mathbb{R}^{n}$ such that, for any hyperplane $H$ separating $\mathbb{R}^{n}$ into halfspaces $H^{-}$and $H^{+}$,
(i) $\operatorname{Pr}_{x \sim \mathcal{D}}\left[x \in H^{+}\right] \geq P$ if and only if $H^{+} \cap K_{P}$ is nonempty, and
(ii) $\operatorname{Pr}_{x \sim \mathcal{D}}\left[x \in H^{+}\right]=P$ if and only if $H^{+} \cap K_{P}$ contains a single point, which is always the centroid of $\mathcal{D}$ restricted to $H$.

The sets $K_{P}$ are called "floating bodies" of $\mathcal{D}$. When $\mathcal{D}$ is a uniform distribution over a compact set $X \subseteq \mathbb{R}^{n}$, a geometric interpretation lends insight. We consider all possible ways to slice off a $P$-fraction of the volume of $X$. After all such slices are made simultaneously, $K_{P}$ is the set that remains. Then the backward direction of statement (i) is immediate, following from the observation that every supporting hyperplane of $K_{P}$ slices off at least a $P$-fraction of the volume of $X$. The nontrivial forward direction is that there are no "redundant" hyperplanes: every plane slicing off at least a $P$-fraction of the volume of $X$ is a supporting hyperplane of $K_{P}$. This part requires the distribution $\mathcal{D}$ to be symmetric and log-concave. It implies that each hyperplane intersects $K_{P}$ at a single point, which must be the center of mass of $X \cap H$ by a straightforward calculus exercise.

The Meyer-Reisner Theorem has received surprisingly little attention in the literature on convex optimization. However, it is intimately tied to convexity, and plays a crucial role in the proofs of our next two theorems.

Theorem 3.9. If the chooser prior $\overline{\mathcal{G}}^{C}$ is symmetric and log-concave, then the set $S_{P}$ of divisions for which the probability the chooser prefers Pile 1 is at most $P$ is a convex set.

Proof. Let $q^{1}, q^{2} \in[-1,1]^{n}$ be two divisions for which the probability the chooser picks Pile 1 is at most $P$. Suppose, toward a contradiction, that for some $0 \leq \alpha \leq 1$, the probability the chooser picks Pile 1 under $q^{3}:=\alpha q^{1}+(1-\alpha) q^{2}$ is $P^{\prime}>P$. In other words, for $i \in\{1,2\}$,

$$
\begin{equation*}
\operatorname{Pr}_{g^{C} \sim \overline{\mathcal{G}}^{C}}\left[q^{i} \cdot g^{C} \geq 0\right] \leq P \tag{7}
\end{equation*}
$$

but

$$
\begin{equation*}
\operatorname{Pr}_{g^{C} \sim \overline{\mathcal{G}}^{C}}\left[q^{3} \cdot g^{C} \geq 0\right]=P^{\prime}>P \tag{8}
\end{equation*}
$$

For each $i \in\{1,2,3\}$, consider the $(n-1)$-dimensional plane

$$
H_{i}:=\left\{g^{C} \in \mathbb{R}^{n} \mid q^{i} \cdot g^{C}=0\right\}
$$

with positive side defined as

$$
H_{i}^{+}:=\left\{g^{C} \in \mathbb{R}^{n} \mid q^{i} \cdot g^{C} \geq 0\right\}
$$

Then Equation (8) says that

$$
\operatorname{Pr}_{g^{C} \sim \overline{\mathcal{G}}^{C}}\left[g^{C} \in H_{3}^{+}\right] .
$$

Applying Theorem 3.8 part (i) (forward direction), we deduce that there is some specific point $g^{C} \in K_{P^{\prime}} \cap H_{3}^{+}$. Since

$$
\alpha\left(q^{1} \cdot g^{C}\right)+(1-\alpha)\left(q^{2} \cdot g^{C}\right)=q^{3} \cdot g^{C} \geq 0
$$

it follows that, for some $i \in\{1,2\}$, we have $q^{i} \cdot g^{C} \geq 0$. Applying Theorem 3.8 part (i) (backward direction) to $H_{i}$, it follows that

$$
\operatorname{Pr}_{g^{C} \sim \overline{\mathcal{G}}^{C}}\left[q^{i} \cdot g^{C} \geq 0\right]=\operatorname{Pr}_{g^{C} \sim \overline{\mathcal{G}}^{C}}\left[g^{C} \in H_{i}^{+}\right] \geq P^{\prime}
$$

contradicting Inequality (7).

Already, this suggests that the technique from Algorithm 1 can be extended beyond normal distributions. When each $\overline{\mathcal{G}}_{i}^{C}$ is Gaussian, the constraint that $q \in S_{P}$ is a convex quadratic constraint. If we only require $\overline{\mathcal{G}}_{i}^{C}$ to be symmetric and log-concave, the constraint is no longer necessarily quadratic, but it is at least convex.

However, convexity alone does not immediately imply we can optimize over $S_{P}$ in polynomial time. For that we must build a separation oracle, which is an algorithm that can efficiently decide whether any given point $q$ is in $S_{P}$, and if it is not, output a hyperplane separating $q$ from $S_{P}$, i.e., a linear constraint that every point in $S_{P}$ satisfies but $q$ does not. It turns out that neither of these tasks is computationally tractable in general. In the case where each $\overline{\mathcal{G}}_{i}^{C}$ is a uniform distribution over an interval, determining whether a given point is in $S_{P}$ is equivalent to computing the volume of an $(n-1)$-dimensional plane cuts out of an $n$-dimensional cube. That well-studied problem is thought to be intractable. ${ }^{3}$ However, our next result shows that we can effectively solve this problem by the means of sophisticated random sampling techniques, using Theorem 3.8 part (ii) to get an (approximate) separating hyperplane. Since the algorithm uses randomness, it is not an FPTAS but rather an FPRAS (fully polynomial-time randomized approximation scheme). An FPRAS is the same as an FPTAS except that it is only required to be correct with some constant probability greater than $\frac{1}{2}$.

Theorem 3.10. An FPRAS exists to compute an approximately optimal division given rational divider values $g_{1}^{D}, g_{2}^{D}, \ldots, g_{n}^{D}$ and chooser priors $\overline{\mathcal{G}}_{i}^{C}=\operatorname{Unif}\left[a_{i}, b_{i}\right]$ encoded by rational numbers $a_{1}, a_{2}, \ldots, a_{n}$ and $b_{1}, b_{2}, \ldots, b_{n}$.

The proof is deferred to Appendix D.

### 3.4 Discrete priors

Using similar ideas we can also efficiently solve the case where the chooser's prior is a discrete distribution when the number of chooser types is small.

Proposition 3.11. There is a function $f$ such that, given a discrete joint distribution $\overline{\mathcal{G}}^{C}$ over the chooser's values (even allowing for correlations across goods), with $n$ goods and $\left|\overline{\mathcal{G}}^{C}\right|=\ell$, it is possible to compute an optimal division in time $f(\ell) \cdot \operatorname{poly}(\mathrm{n})$.

Algorithm 2, presented in Appendix E, solves this problem by optimizing a similar ensemble of linear programs and taking the best solution. Unlike our other algorithms, this returns an exact (rather than approximate) solution and works even when there are correlations among the chooser's values for the various goods.

## 4 Diversification: Why and how

To maximize his expected utility, the divider must balance two objectives when he allocates goods to the piles: maximizing the returns from the more-desirable pile 1 , and reducing the risk that the chooser picks pile 1. The divider trades off between these objectives by transferring just the right amount of value into pile 2 . One might naturally expect that the divider's strategic approach from the case of known chooser preferences would still be optimal. Namely, the divider could start with all goods in pile 1 and transfer goods into pile 2 in order of ascending critical ratios, thus creeping

[^2]along the risk-return frontier to find the optimal utility. In the end, at most one good would need to be divided.

However, another strategic factor-diversification-can reduce the risk to the divider that the chooser will pick pile 1, and thus enable the divider to do better. In the investment context, investors know they can push the entire risk-return frontier outwards by investing in many assets. This analogy applies quite well to our setting: given that the chooser is more likely to pick pile 2 (which is always the case by Lemma 3.2), when the expected difference in the chooser's values for the two piles is fixed, she is more likely to pick pile 2 if the variability of the difference is lower. Thus, the optimizing divider should not merely transfer a sufficient amount of value into pile 2 . He should also reduce variance in this value by dividing multiple goods between the two piles, thereby diversifying the piles to reduce risk. In this section, we analyze how this incentive affects the divider's optimization problem.

If goods are indivisible, diversification can still be achieved by using lotteries, whereby the divider does not merely divide the goods into two piles, but commits to random allocations of certain goods after the chooser has made her choice. Consider goods 1 and 2 (of many others), for which the chooser's values are equally likely to be 0 or 1 . If good 1 is put in pile 1 and good 2 in pile 2, then there is a $25 \%$ chance that the chooser's value of those two assets will be greater in pile 1 than in pile 2 . By contrast, if each good is put in pile 1 with a $50 \%$ chance, and pile 2 with a $50 \%$ chance, then the chooser-before the lotteries have been conducted-will always value the probabilistic assets in the piles equally. We emphasize here that the "diversification" imperative under examination is not aimed at reducing the risk in the lotteries for each good, but instead at reducing the risk that the chooser picks pile 1 before the lotteries are resolved. If the divider is risk-averse, then the incentive to diversify over risk in the chooser's action can be at odds with an inherent aversion to using lotteries. This is a subtle issue we address briefly in Section 4.3, when we discuss risk-aversion given both divisible goods and indivisible goods that can be divided via lotteries.

### 4.1 Which goods get divided?

We know that the divider can find the optimal division of goods in the normal and discrete cases using Algorithm 1 and Algorithm 2 respectively. Here we provide a qualitative explanation of how goods are optimally allocated between the two piles.

We begin by observing that it may be optimal for the divider to split all but one good, as Figure 2 illustrates. The most valuable good-worth about twice as much as any of the others-is placed entirely in pile 1 , while the other goods are all placed mostly in pile 2 but split between the two piles in order to reduce variance, pushing $P$ extremely close to zero. This optimal division was computed using Algorithm 1. ${ }^{4}$

This result starkly contrasts with the case of complete information. As we discussed in Section 1.2 , a divider with perfect information never needs to split more than one good between the piles.

[^3]The Optimal Division of Six Goods with Five Goods Split


Figure 2: An example with six goods where it is optimal for the divider to split five of the goods between the two piles. Here $g_{1}^{D}=9.8, g_{2}^{D}=9.9, g_{3}^{D}=10, g_{4}^{D}=10.1, g_{5}^{D}=10.2, g_{6}^{D}=15$, and the chooser's prior for each good is $\overline{\mathcal{G}}_{i}^{C}=\mathcal{N}(10,1)$. The optimal value of $P$ is 0.078 .

Figure 3 interpolates between these two cases by decreasing the variance from one down to zero. One by one, the high value goods are placed entirely in pile 1 and the low value goods are placed entirely in pile 2. Eventually, as the variance approaches zero, we arrive at the certainty case, where there is no longer any diversification.

Although the first five goods are diversified in this example, the ones the divider values more tilt more toward pile 1 than do the ones the divider values less. This suggests an intuitive rule by which the divider could be guided in constructing his optimal division. Recall that, in the absence of uncertainty, the fraction $p_{i}$ of good $i$ placed in pile 1 is monotonically increasing in the critical ratio of good $i .{ }^{5}$ Does this monotonicity result survive in the presence of uncertainty?

While it seems intuitive that prioritization by critical ratios would carry over to the uncertainty case, it will not if the priors have different variances. Even if a good has a very large or very small critical ratio, its variance may be so large to make it too risky to place it mostly in one pile. Thus, when variances differ among goods, monotone divisions may be suboptimal. Surprisingly, we observe that monotonicity may be lost even when goods have normal priors with identical variances.

Figure 4 shows the optimal division in a simple example with three normally-distributed goods, again computed using Algorithm 1. Even though good 2 has a higher critical ratio and the same variance as good 1 , the optimal division sets $p_{2}<p_{1}$. The primary reason is that the value of good 2 to the chooser, in absolute terms (rather than relative to the divider's value), is so large that,

[^4]

Figure 3: Optimal divisions as a function of the variance $s$ of the chooser's value for each good. This is the same setup as in Figure 2, but instead of chooser values being drawn from $\mathcal{N}(10,1)$, they are drawn from $\mathcal{N}(10, s)$, for $s$ starting at one and approaching zero. Thus the first (dark blue) bars are the same division as before.
in order to ensure that the value of $P$ stays low, raising $p_{2}$ would require lowering $p_{1}$ by such a large amount that the divider would be worse off. Note that we do require a high variance of 25 for each good to see this effect. As we decrease the variance to zero, we see convergence to the case of known preferences, with the division becoming monotone in the critical ratios starting around a variance of 5 .

This example is surprising given the small amount of variance relative to the mean of each good. While in the polar case of zero variance, monotonicity in the critical ratios is optimal, a sliver of uncertainty breaks this result.

A further departure from the case of certain preferences is that the divider can achieve a higher-than-proportional utility even when all critical ratios are the same - in fact, even when the chooser's prior distributions are identical and the divider values every good identically. Clearly, such a division is not feasible with normal priors by Proposition 3.3. With two-point (or multipoint) discrete priors, despite such identical values across goods, it is feasible for the divider to beat his proportionality guarantee.

Proposition 4.1. There exists a number of goods $n$ and a discrete distribution $\mathcal{D}$ supported on two positive values such that, even if, for each good $i \in[n], g_{i}^{D}=1$ and $\overline{\mathcal{G}}_{i}^{C}=\mathcal{D}$, there exists a division of goods yielding utility higher than the divider's proportionality guarantee.

Proof. Let $n=5$ and let $\mathcal{D}$ be the following distribution: value 0.01 with probability 0.6 and value


Figure 4: An instance with three goods where, despite all prior variances being the same, the $p_{i}$ are not monotone in the critical ratios. Here $g_{1}^{D}=1, g_{2}^{D}=2, g_{3}^{D}=3$, and corresponding chooser priors are $\overline{\mathcal{G}}_{1}^{C}=\mathcal{N}(100,25), \overline{\mathcal{G}}_{2}^{C}=\mathcal{N}(198,25), \overline{\mathcal{G}}_{3}^{C}=\mathcal{N}(100,25)$. The optimal value of $P$ is 0.005 .

1 with probability 0.4 . Consider the division

$$
p=(1,0.4,0.4,0.4,0.4) .
$$

Observe that if the chooser values the first good at 0.01 and at least one of the other goods at 1 , then she prefers pile 2 . This happens with probability

$$
0.6 \cdot\left(1-0.6^{4}\right)=0.553344>\frac{1}{2} .
$$

Since the divider values pile 1 more than pile 2, he achieves a higher utility than his proportionality guarantee; specifically, his expected utility is $2.510068>2.5$.

Finally, we note that there is a limit to the need to diversify through lotteries. It is never strictly optimal to split all goods between the two piles. The proof is straightforward, and is included in Appendix F for completeness.

Lemma 4.2. If the divider can achieve a utility exceeding his proportionality guarantee, he will always leave one good entirely in one of the two piles.

We also remark that, even though all examples above place an undivided good optimally in pile 1 , sometimes it is optimal to place a sole undivided good in pile 2 .

### 4.2 The effects of correlation between the players' values

Thus far we have only considered settings where the divider's prior over the chooser's values is specified arbitrarily and is independent of the divider's values. In this section we briefly study optimal divisions embedded in a larger, realistic context where there is a joint prior over both players' values. Such a situation would arise if there were a public value component to the player's valuations, as say for market value when jewelry pieces or real estate assets are being divided. The divider uses his own values to update his beliefs about the chooser's values.

Specifically, consider a scenario where both players' values follow a Gaussian marginal distribution $\mathcal{G}_{i}^{D}=\mathcal{G}_{i}^{C}=\mathcal{N}(10,1)$ for each good $i$. However, the values for each good are correlated. For each good $i$, there is a common public value $g_{i}$ that is distributed according to $\mathcal{N}(10, t)$, for a parameter $t \in[0,1]$. Additionally, each player has a private value $\varepsilon_{i}^{D}, \varepsilon_{i}^{C}$ distributed independently according to $\mathcal{N}(0,1-t)$. The values for good $i$ are $g_{i}^{D}:=g_{i}+\varepsilon_{i}^{D}$ and $g_{i}^{C}:=g_{i}+\varepsilon_{i}^{C}$.

Optimal Divisions Under Stronger Value Correlations


Figure 5: Optimal divisions as a function of the correlation parameter $t$. As in Figure 2, each value is is drawn from $\mathcal{N}(10,1)$, and the divider's values happen to be $g^{D}=(9.8,9.9,10,10.1,10.2,15)$. The difference is that now, for $t>0$, the divider infers that the chooser has similar values to his own. The first (dark blue) bars corresponding to $t=0$ are the same division as before. As $t$ approaches 1 , the optimal division changes little from the final (red) bars at $t=7 / 8$.

Figure 5 shows the optimal divisions for increasing values of $t$ in the same setting as before, where the divider has one good at a very high value of 15 and the other goods' values are close to the mean of 10 . As is evident, the divider's general strategy does not change in any significant qualitative way. The high-value good is still placed entirely in pile 1 , and the divider still diversifies among the other goods. Perhaps surprisingly, as $t$ increases toward 1 , the divider actually diversifies more, with the gaps between the fractions in pile 1 among the first 5 goods shrinking.

One possible explanation for this trend is as follows. As $t$ increases, the divider increasingly expects the chooser to have similar values, as most of the variation can be attributed to the public value. It thus becomes more important for the divider to divide "evenly." Otherwise, if the divider concentrates too much value in pile 1 , then because the chooser has similar values, the chooser will pick pile 1. But then, as the piles start to have similar values to the divider, the consequences of the chooser picking pile 1 become less dire. So the divider will optimally choose a division with a higher value of $P$. (Indeed, this is the case in this example. In going from $t=0$ to $t=1 / 2$, the value of $P$ doubles, then continues to approach $1 / 2$ as $t$ approaches 1.) In this regime, diversification becomes more important, as decreasing variability in the values of the two piles can have an even greater effect on improving the chances the chooser picks pile 2 .

### 4.3 The effects of risk aversion

It is optimal for the divider to increase his expected value by taking a risk on the value he receives; how does his strategy adjust if he is averse to risk? Suppose the divider is an expected utility maximizer with utility $u^{D}=f\left(v^{D}\right)$, where $f$ is an increasing concave function and $v^{D}$ is the total value of all goods the divider receives.

Thus far, we have not distinguished between deterministic divisions of divisible goods versus randomized divisions of indivisible goods. For example, if $p_{i}=\frac{1}{2}$, then it could mean that good $i$ is literally split between the piles into two equal pieces, or that it will be randomly allocated, in whole, to one pile or the other after the chooser picks a pile. In the two cases the two players' incentives are the same. However, if the divider is risk-averse, a lottery that is resolved after piles are selected will impose unwanted risk on him. In this section, we assume that the goods are divisible and note where results would be different if the goods were indivisible and allocated by lottery.

We identify two main effects of a divider's risk-aversion. First, it decreases the probability $P$ that the chooser picks pile 1, as the following theorem shows.

Theorem 4.3. Assume all values of goods are always nonnegative. Fix divider values and let $f$ be a concave utility function. The probability that the chooser picks pile 1 under any optimal division by a risk-neutral divider is weakly greater than the probability the chooser picks pile 1 under any optimal division by a divider with utility function $f$.

Proof. Let $p$ be the optimal division by a risk-neutral divider, and let $p^{\prime}$ be the optimal division by a risk-averse divider with utility function $f$. Let $T$ be the total divider value of all goods. Let $v, P$, and $v^{\prime}, P^{\prime}$ denote the divider's value for pile 1 and the probability the chooser picks pile 1 according to $p$ and $p^{\prime}$, respectively. Suppose toward a contradiction that $P<P^{\prime}$. Then we must have $v<v^{\prime}$, for otherwise $p$ would be a strictly better division than $p^{\prime}$ for the risk-averse divider, as it would simultaneously yield a higher value in pile 1 and higher probability of the chooser picking pile 2. Also, since the risk-neutral divider prefers $q$ to $q^{\prime}$, it must be that

$$
\left(1-P^{\prime}\right) v^{\prime}+P^{\prime}\left(T-v^{\prime}\right)<(1-P) v+P(T-v) .
$$

Thus, values are ordered exactly as shown in Figure 6, where the $x$-coordinate of points $A$ and $C$ are the respective expected divider values of $p$ and $p^{\prime}$. If the solid blue curve is the utility function $f$, then the $y$-coordinates of points $A$ and $C$ are the expected utilities of the risk-averse divider using $p$ and $p^{\prime}$, respectively. It follows from monotonicity and convexity that $A$ must have a higher $y$-coordinate than $B$, which must have a higher $y$-coordinate than $C$. This contradicts our assumption that the risk-averse divider prefers $p^{\prime}$.

## Divisible Goods Case: Risk-Averse Divider Selects a Smaller $P$



Figure 6: Illustration accompanying the proof of Theorem 4.3.

We may observe this effect empirically by assuming the utility function $f(x)=\sqrt{x}$. This turns the convex program $\mathcal{C}_{P}$ from Algorithm 1, which maximizes divider utility with respect to a fixed value of $P$, into a non-convex program. However, it can be written with only two non-convex constraints, so it is still practically feasible to solve it exactly. Figure 7 plots the optimal expected divider utility given any value of $P$ for a divider with utility function $f(x)=\sqrt{x}$. The divider's values and chooser's priors are the same as in Figure 1. Comparing the two figures, one can see that risk aversion reduces the optimal utility for larger values of $P$, making smaller values of $P$ more attractive.

Theorem 4.3 does not hold when goods are indivisible and allocated randomly. Consider a scenario with only $n=2$ goods. Good 1 deterministically has value 4 for both players. Good 2 is worth 16 to the divider and either 1 or 12 to the chooser, each with probability $\frac{1}{2}$. A risk-neutral divider will put $\frac{2}{3}$ of good 2 in pile 1 and everything else in pile 2 , ensuring that the chooser always picks pile 2 (i.e., $P=0$ ). If the divider has utility function $f(x)=\sqrt{x}$, then assuming " $\frac{2}{3}$ of good 2 " is a lottery over good 2 , this strategy yields expected utility $\frac{2}{3} \sqrt{16}=\frac{8}{3}$. However, by just putting the goods in separate piles and not using lotteries, the chooser will pick pile 2 with probability $P=\frac{1}{2}$, yielding expected divider utility

$$
\frac{1}{2} \sqrt{4}+\frac{1}{2} \sqrt{16}=3>\frac{8}{3}
$$

We verified computationally that no division in which $P=0$ yields expected divider utility greater than $\frac{8}{3}$.


Figure 7: The same plot of expected divider utility versus the probability that the chooser picks pile 1 as in Figure 1, but with a risk averse divider. This example uses the same divider values and chooser priors as in Figure 1. Again, the globally optimal value of $P$ is indicated as $P^{*}$.

The second major effect of divider risk aversion is to increase the amount of diversification. Consider the setting with 40 goods with both divider and chooser values drawn i.i.d. from $\mathcal{N}(1,0.04)$. Figure 8 compares the optimal divisions by a risk-neutral divider and a risk-averse divider using the utility function $f(x)=\sqrt{x-5}$ with values for the 10 goods drawn independently from $\mathcal{N}(1,0.04)$. As one can see, risk-aversion leads to more goods being split between the two piles, and they are generally split more equally. Still, one good is always left undivided, as Lemma 4.2 holds with the same proof.

## 5 Welfare of the players

We now turn to analyze the expected welfare of the players. Knowing their expected welfare is important if our concern is some conception of fairness. It is also critical in enabling a player to decide when to try to play divider, and when chooser. The divide and choose game is explicitly asymmetric, and despite its minimal axiomatic guarantees (e.g., envy-freeness), one player might end up better off than the other merely due to this asymmetry. In this section we identify conditions under which one role is preferable to the other.

### 5.1 Number of goods

The cake-cutting literature suggests that the chooser is better off for the unknown-preferences case, as the divider is compelled to divide goods roughly evenly, while the chooser can get a more favorable


Figure 8: Comparison of the divider's optimal divisions under risk-neutrality (left, dark blue bars) and risk-aversion (right, light red bars). The chooser's prior for each good $i$ is $\mathcal{G}_{i}^{C}:=\mathcal{N}(1,0.04)$, and the divider's values were also independently sampled from $\mathcal{G}_{i}^{D}:=\mathcal{N}(1,0.04)$. Each bar represents a value of $p_{i}$, and is horizontally located at the divider's value $g_{i}^{D}$ (indicated by the dashed black lines).
outcome because she will typically value one pile more than the other. However, if the divider has strong knowledge of the chooser's preferences, then the divider can exploit this knowledge [27, 9]. Indeed, as Nicoló and Yu [23] note in the context of cake-cutting, "The divide and choose rule leads to a no-envy outcome but the rule itself is not envy free: the chooser envies the role of the divider." Thus, the best generalization one could hope to make is that the relative utilities of the two players depend on the amount of uncertainty faced by the divider. Hence, the greater one's own and one's counterpart's knowledge, the greater the benefit of playing the divider. Weak knowledge, on either side, favors being the chooser.

One natural example of this phenomenon is when values for all $n$ goods are drawn i.i.d. from the same distribution. If $n$ is small, there is significant uncertainty in how the chooser will value the piles. Moreover, the divider cannot count on receiving what he places in pile 1. Consequently, when goods are few, the chooser has the advantage.

In contrast, when $n$ is large, the uncertainty in the value of a pile shrinks relative to its mean value. In this situation, the divider can cluster his high-value goods into pile 1, and expect to receive that pile with high probability. Thus, the divider is favored. These observations lead to the final major result in this paper: For plausible distributions of players' values, the chooser is favored when $n$ is small; the divider is favored when $n$ is large. In the latter case, this is true even when values are correlated (positively or negatively) between the players.

Theorem 5.1. Let $\mathcal{D}$ be a probability distribution over $\mathbb{R}_{>0} \times \mathbb{R}_{>0}$. Suppose the values for each good $i$ are drawn from $\mathcal{D}$, independently across $i$. Then, as long as the critical ratio of each good is not always the same value with probability one, the following hold in any equilibrium.
(i) If the two components of $\mathcal{D}$ are independent, then for $n=2$ it is strictly better to be the chooser:

$$
\underset{\left(g^{D}, g^{C}\right) \sim \mathcal{D}^{n}}{\mathbb{E}}\left[u^{D}\right]<\underset{\left(g^{C}, g^{D}\right) \sim \mathcal{D}^{n}}{\mathbb{E}}\left[u^{C}\right]
$$

(ii) If, for all $\varepsilon>0$, for all sufficiently large $n$,

$$
\begin{equation*}
\operatorname{Pr}_{\left(g^{1}, g^{2}\right) \sim \mathcal{D}}\left[g^{2} \leq 8 n \exp \left(-2 \varepsilon^{2} n\right)\right]<\frac{\varepsilon}{n}, \tag{9}
\end{equation*}
$$

then for all sufficiently large $n$, it is strictly better to be the divider:

$$
\underset{\left(g^{D}, g^{C}\right) \sim \mathcal{D}^{n}}{\mathbb{E}}\left[u^{D}\right]>\underset{\left(g^{C}, g^{D}\right) \sim \mathcal{D}^{n}}{\mathbb{E}}\left[u^{C}\right]
$$

As an illustration of Theorem 5.1, take $\mathcal{D}$ to be the uniform distribution on $[0,1]^{2}$ : every player's value for every good is drawn uniformly from the unit interval. This easily satisfies all the assumptions of Theorem 5.1. Table 1 lists the ex-ante expected utilities for each player with two goods and with many goods. With two goods, the chooser is about $50 \%$ better off; with many goods, the divider is $50 \%$ better off.

Why does this happen? In the two-good case, the key observation is that the divider will always fractionally allocate his favorite good, mostly in pile 1 , placing his less favorite good entirely in pile 2. Thus, when the chooser's favorite good is different from the divider's, she always does better, since she receives all of the good rather than part of it. When they prefer the same good, one can likewise see that the chooser is still better off via a simple case analysis over the relative fractions of the favorite good that each player would have kept in pile 1 had they been the divider.

For large $n$, relying on the law of large numbers, the divider could be confident that if he placed slightly more than $50 \%$ of the goods in pile 2 , the chooser would pick pile 2 with an arbitrarily high probability. For example, when $n=100$, on average, the optimizing divider puts his 45 highestvalued goods in pile 1, and his 55 lowest-valued goods in pile 2 . The chooser picks pile 1 with probability $P=0.04$. In the limit as $n \rightarrow \infty, P$ tends to zero, so the divider's ex-ante utility is the sum of his top-half most-valued goods. On the other hand, for very large $n$, by the law of large numbers, the chooser can expect only her proportionality guarantee utility from each pile, i.e., the sum of a random subset of half of her values. This is essentially the whole argument of the proof, though establishing the trends in full generality, even when values are correlated, requires carefully defining and analyzing the optimal divisions.

| Equilibrium utility per $U[0,1]$ good | $n=2$ | $n \rightarrow \infty$ |
| ---: | :---: | :---: |
| Proportionality guarantee | 0.25 | 0.25 |
| Divider | $\frac{19}{72}=0.264$ | 0.375 |
| Chooser | $\frac{1031}{2160}+\frac{\ln (3 / 4)}{3}=0.381$ | 0.25 |

Table 1: Expected utilities of the two players, normalized by the number of goods, when the value each player has for each good is drawn i.i.d. from the uniform distribution on $[0,1]$. The $n=2$ column was computed using Mathematica. The $n \rightarrow \infty$ column was computed using the proof of Theorem 5.1 (ii).

In Appendix G we remark on several ways in which the assumptions are necessary and conclusions as strong as possible. Below, we give a complete proof for part (i) because it is short and intuitive. The proof of part (ii), which involves tedious approximations and concentration bounds, is deferred to Appendix H.

Proof of Theorem 5.1 part (i). By Lemma 4.2, the divider will always leave one good undivided. Since the chooser's values for both goods are drawn independently from the same distribution $\mathcal{D}$, which is supported only over positive values, the chooser is more likely to pick the pile with the undivided good. Therefore, it must be that the divider's least-preferred good is the one that remains undivided, for otherwise he could have a higher expected utility switching the roles of the two goods. In other words, for some function $f$ (that depends on the fixed distribution $\mathcal{D}$ ), the optimal division is always as follows: given values $g_{1}^{D}$ and $g_{2}^{D}$, the divider places a $p^{D}:=f\left(\left\{g_{1}^{D}, g_{2}^{D}\right\}\right)>\frac{1}{2}$ fraction of good $\arg \max _{i}\left(g_{i}^{D}\right)$ in pile 1, with the rest of that good and all of the other good in pile 2. Analogously, let $p^{C}:=f\left(\left\{g_{1}^{C}, g_{2}^{C}\right\}\right)$, i.e., the amount of the chooser's preferred good that would have gone into pile 1 had the chooser been the divider.

To compare the two players' ex-ante expected utilities, we fix realizations of $g_{1}^{D}, g_{2}^{D}, g_{1}^{C}$, and $g_{2}^{C}$, and compare the chooser's actual utility with the hypothetical utility had the roles been reversed. There are three possible cases to consider, depending on the realizations of $g_{1}^{D}, g_{2}^{D}$, $g_{1}^{C}$, and $g_{2}^{C}$.

Case 1: The two players weakly prefer different goods. In this case, the chooser will receive all of her favorite good in pile 2 . Had she instead been the divider, she would have only received a $p^{C}$-fraction of her favorite good. Since $p^{C} \leq 1$, the chooser is weakly better off.

Case 2: The two players strictly prefer the same good, and $p^{D} \geq p^{C}$. In this case, the chooser receives a $p^{D}$-fraction of her favorite good in pile 1. As in Case 1, had she instead been the divider, then she would have only received a $p^{C}$-fraction of her favorite good. Since $p^{C} \leq p^{D}$, the chooser is weakly better off.

Case 3: The two players strictly prefer the same good, and $p^{D}<p^{C}$. In this case, had the chooser been the divider, the divider would have opted for the chooser's preferred pile 1 , since we know the divider would have weakly preferred a $p^{D}$-fraction of his favorite good to everything else, so he must have strictly preferred a $p^{C}$-fraction of his favorite good. Thus, if roles had been reversed, the chooser would have received her least-preferred pile, obtaining utility at most her proportionality guarantee. Since the chooser always can get at least this utility, the chooser is weakly better off being the chooser.

In short, we have shown that the chooser is weakly better off in all cases. Furthermore, the chooser is strictly better off in Case 1 whenever $p^{C}<1$. If this never occurs, that means that the divider always puts the goods entirely into different piles, in which case the divider receives utility equal to his proportionality guarantee and the chooser exceeds her proportionality guarantee (since the assumptions on $\mathcal{D}$ imply the chooser's values for the two goods are not always the same), so we are done. Assuming that $p^{C}<1$ with nonzero probability, we further observe that, conditioning on $p^{C}<1$, the players will prefer different goods (or be indifferent) at least half of the time. Thus, we see that the chooser is strictly better off with nonzero probability, so she is strictly better off overall.

Theorem 5.1 appears to hold for normal priors with a positive mean as well (even though they may take unbounded and negative values). For $n=2$, the chooser is better off, while for large $n$, the divider is better off. At what value of $n$ does the crossover occur? Figure 9 plots estimated utilities per good for an empirical experiment where all values are drawn from $\mathcal{N}(1,0.04)$. The crossover appears to be around 15 goods.


Figure 9: Estimated utilities from repeated trials of the following experiment with independent values. We first draw divider values $g_{1}^{D}, g_{2}^{D}, \ldots, g_{n}^{D}$ i.i.d. from $\mathcal{N}(1,0.04)$, then compute the optimal division $p$ with respect to those values, with $\mathcal{G}_{i}^{C}=\mathcal{N}(1,0.04)$ for each good $i$ as well. We compute the divider's and chooser's expected utility exactly with respect to $p$ and $g_{1}^{D}, g_{2}^{D}, \ldots, g_{n}^{D}$, where the expectation is taken over the unknown values $g_{1}^{C}, g_{2}^{C}, \ldots, g_{n}^{C}$. All utilities are averaged, then normalized by dividing by the number of goods $n$.

### 5.2 Heterogeneity of values

Thus far we have been measuring and comparing the utilities afforded by the two roles ex-ante. However, there is more to the story: the realizations of the values will generally affect which role is more desirable. Figure 10 plots random samples of the values of 13 goods drawn from the same distribution, $\mathcal{N}(1,0.04)$, as in Figure 9. We chose the number of goods to be 13 since that is near the crossover point in Figure 9, where the two roles have similar ex-ante utilities. ${ }^{6}$ As one can see, when values vary more widely it is better to be the divider. This is because the divider's optimal strategy is to place all goods that he values highly in pile 1 ; he cares much less about the low-value goods. On the other hand, when values are more concentrated around the mean, it is better to be the chooser. In such a scenario, the divider's strategy will not yield a substantially high payoff. Instead, the chooser benefits simply from the fact that she will probably be able to take a larger and more valuable pile under the divider's optimal strategy.

[^5]

Figure 10: 200 random samples of values for 13 goods drawn i.i.d. from $\mathcal{N}(1,0.04)$, colored by the role that is better off given such values. Each row of points represents a separate draw, and the draws are sorted from bottom to top in increasing sample deviation (sum of absolute differences between the value of good $i$ and the mean of 1 , across all 13 goods). Divider utilities are computed as the optimal utilities from Algorithm 1. Chooser utilities are estimated by averaging utilities with respect to a fixed ensemble of 4,000 optimal divisions under random divider values drawn from $\mathcal{N}(1,0.04)$.

## 6 Conclusions

Divide-and-choose methods are widely employed by business partners going separate ways, inheritors dividing assets, and divorcing couples. They also lie at the heart of an array of managerial decisions.

The divide-and-choose game is not unfamiliar to economics and computer science. However, typical formulations - such as the famed cake-cutting model - assume the players' preferences are known, or at the very least can be elicited truthfully. In the real world, by contrast, a divider will merely have subjective beliefs about the chooser's preferences. This analysis pioneers the study of optimal strategies in divide-and-choose when one's counterpart's preferences are unknown.

A major finding is that the known and unknown preference situations lead to qualitatively different strategies. Moreover, a number of the results for the unknown preferences case are surprising. For example, with known preferences, the ratio of values for the two players, what we label the critical ratio, proves critical. All the goods in one pile have a higher ratio than any of the goods in the other. In addition, at most one good is divided between the piles, with division via a lottery if that good is physically indivisible. With unknown preferences, by contrast, assignment by critical ratios may be violated. That is, fractions allocated to pile 1 may be non-monotonic in the ratios.

Moreover, it might be optimal to fractionally allocate up to $n-1$ out of $n$ goods.
Uncertainty on preferences makes diversification a vital concern. It competes with efficient division as an objective for the divider. The analysis proceeds with a pile 1 , preferred by the divider, with its complement denoted pile 2. The divider is eager to get the chooser to pick the latter. This is fostered in part by assigning greater chooser expected value to pile 2 . This process is limited because the total value in pile 1 is being diminished, and that is the pile the divider hopes to get and will get most of the time. Holding the disparity in expected values fixed, the chooser is more likely to opt for pile 1 the greater is variability in actual values. Diversification-achieved by dividing multiple goods between the two piles - significantly reduces the variability in the chooser's valuation of piles. Hence, it reduces the likelihood that the divider ends up with pile 2 .

Computing the divider's optimal division is easy in the known-preferences case. It is surprisingly difficult with preferences unknown. We already mentioned the troubling lack of monotonicity in critical ratios. Worse still, multiple local optima are commonly encountered. Despite these challenges, we are pleased to identify algorithms that compute divisions for normal and uniform priors in polynomial time, approximating the optimal utility to within arbitrary precision. We suspect that exactly computing optimal divisions in either of these settings is NP-Hard. Characterizing the precise computational complexity is an interesting open question.

We investigated some general effects of divider risk aversion. While this makes the divider's optimization task more computationally difficult, our results do shed some light on how to actually find optimal divisions. Chooser risk aversion, not addressed here at all, merits future study. If goods are divisible, chooser risk aversion is irrelevant. With indivisible goods only divided via lotteries, it could be consequential. One intriguing question is whether correlating lotteries across multiple goods can increase the divider's expected utility.

Many of our results have direct implications for real-world divide-and-choose situations, implying that the models from which they flow represent stylized facts. For example, a typical buy-sell arrangement allows either player to grab the yet unclaimed divider's role. The relative attraction of the divider's role is greater the more refined one's assessment of the other player's preferences, the better informed the other player about your preferences, and the greater the number of goods. Both the divider and the chooser roles gain in value the more extreme one's preferences relative to expectation; the parameters of the specific situation will determine which role is preferred.

Insights from divide-and-choose apply in a range of contexts that are not specifically divide-andchoose. A closely related setting occurs, for example, when one party makes a take-it-or-leave-it offer, such as a job offer, to another party. The recipient can accept the offer or stick with the status quo. Such offers commonly include divisible goods. Thus, the offer of a university to a job candidate might specify a salary and a teaching load, where there is a permissible range for each attribute.

There is an indirect equivalent to our formulation. An accepted offer, the university's preferred outcome, is the equivalent of pile 1. A rejected offer is the status quo. (No resources change hands given a rejection.) The chooser accepts when he prefers the offer to the status quo. In the traditional divide-and-choose situation, the divider is uncertain about the chooser's valuation of the two piles. In a take-it-or-leave it offer, she is uncertain about the chooser's valuation of the offer relative to the status quo. On a technical level, these two games are slightly different, where the latter game essentially involves the same optimization problem except that the $q_{i}$ variables must be nonnegative and $P$ can be greater than $\frac{1}{2}$. Thus we would expect our analyses about critical ratios and diversification to still apply.

Posit that the university and the candidate have linear and additive preferences over the two attributes. However, the university has only priors on the candidate's weights on attributes and valuation of the status quo. It would like to increase the probability of acceptance, i.e., that the
offer lies above the candidate's unknown reservation value, but also wants to increase its value from an accepted offer. An offer using interior values for some attributes - a direct analogue to fractional allocations in our model-will often be optimal. Interior values provide diversification; they increase the probability of acceptance for a given value of an accepted offer to the university.

As a final example, many bargaining contexts involve reciprocal concessions on different issues over time as a means to build confidence. Thus, a labor union might initially give up on its four-day-a-week demand in exchange for the corporation's willingness to fund pensions more generously. When offering such preliminary trades, each party will reason about critical ratios to decide the optimal offer or counteroffer. But it will also be thinking ahead as to what it will have to give up on the big issue, say salary, to get an acceptance. While our model is formally situated as a one-shot interaction, we believe our insights extend qualitatively to these more complex negotiations: at every step on the path to a deal, the parties must think about both critical ratios and effective diversification.

There are intriguing areas for future work, only touched on in this paper, that will require a blend of analytic advances and attention to real-world features. They include player risk aversion, ex-post renegotiation of agreed-upon allocations, and mechanisms-such as alternating choicesthat mix responsibilities for playing divider and chooser.

Of course, a salient feature of the real world is that players are not perfectly rational. An interesting theoretical question for future work would be to investigate an alternative equilibrium concept where players make mistakes with small probability. Also, an experimental study is currently underway to see how dividers and choosers actually play in a game closely resembling our formulation.

Allocation systems, in a range of areas from divide-and-choose to the market, involve players acting strategically to maximize their take given what they know of the other player's valuation. This paper hopes to have shown that, in the allocation of analytic attention, the art of actually playing the divide-and-choose game has received a smaller pile than it deserves.

## Acknowledgments

We are deeply grateful to Jack Stade, who suggested the proof approach for Theorem 3.9; to Pranay Gorantla, who found and alerted us to the Meyer-Reisner Theorem (Theorem 3.8); and to Yannai Gonczarowski, for his careful reading of an earlier draft with extremely helpful feedback.

We also thank Justin Chan for numerous suggestions that improved readability, as well as our anonymous reviewers at both Management Science and EC 2023. Finally, we thank Paul Gölz, Gregory Kehne, Prashanth Amireddy, Elizabeth Pratt, and Santhoshini Velusamy, who also consulted on Theorem 3.9, and David Eppstein, who consulted on Theorem 3.10.

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE1745303, and by the Mossavar-Rahmani Center for Business and Government, Harvard University. Any opinions, findings, and conclusions or recommendations expressed in this material reflect the views of the authors and not necessarily those of the National Science Foundation or the Mossavar-Rahmani Center.

## References

[1] Georgios Amanatidis, Georgios Birmpas, Federico Fusco, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. Allocating indivisible goods to strategic agents: Pure nash equilibria
and fairness. In Web and Internet Economics - 17th International Conference, WINE 2021, Potsdam, Germany, December 14-17, 2021, Proceedings, pages 149-166, 2021.
[2] Haris Aziz and Simon Mackenzie. A bounded and envy-free cake cutting algorithm. Communications of the ACM, 63(4):119-126, 2020.
[3] Siddharth Barman, Ganesh Ghalme, Shweta Jain, Pooja Kulkarni, and Shivika Narang. Fair division of indivisible goods among strategic agents. In Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19, Montreal, QC, Canada, May 13-17, 2019, pages 1811-1813, 2019.
[4] Xiaohui Bei, Ning Chen, Guangda Huzhang, Biaoshuai Tao, and Jiajun Wu. Cake cutting: Envy and truth. In Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017, pages 3625-3631, 2017.
[5] Xiaohui Bei, Guangda Huzhang, and Warut Suksompong. Truthful fair division without free disposal. Soc. Choice Welf., 55(3):523-545, 2020.
[6] Anna Bogomolnaia, Hervé Moulin, and Fedor Sandomirskiy. On the fair division of a random object. Management Science, 68(2):1174-1194, 2022.
[7] Steven J Brams and Samuel Merrill, III. Equilibrium strategies for final-offer arbitration: There is no median convergence. Management Science, 29(8):927-941, 1983.
[8] Steven J. Brams and Alan D. Taylor. An envy-free cake division protocol. The American Mathematical Monthly, 102(1):9-18, 1995.
[9] Simina Brânzei, Ioannis Caragiannis, David Kurokawa, and Ariel D. Procaccia. An algorithmic framework for strategic fair division. In Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA, pages 418-424, 2016.
[10] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. Journal of Political Economy, 119(6):1061-1103, 2011.
[11] Yiling Chen, John K. Lai, David C. Parkes, and Ariel D. Procaccia. Truth, justice, and cake cutting. Games and Economic Behavior, 77(1):284-297, 2013.
[12] Yun Kuen Cheung. Better strategyproof mechanisms without payments or prior - an analytic approach. In Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016, pages 194-200, 2016.
[13] Richard Cole, Vasilis Gkatzelis, and Gagan Goel. Mechanism design for fair division: allocating divisible items without payments. In Proceedings of the fourteenth ACM Conference on Electronic Commerce, EC 2013, Philadelphia, PA, USA, June 16-20, 2013, pages 251-268, 2013.
[14] Vincent Conitzer, Rupert Freeman, and Nisarg Shah. Fair public decision making. In Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017, pages 629-646, 2017.
[15] Payam Delgosha and Amin Gohari. Information theoretic cutting of a cake. IEEE Transactions on Information Theory, 63(11):6950-6978, 2017.
[16] Martin E. Dyer and Alan M. Frieze. On the complexity of computing the volume of a polyhedron. SIAM Journal on Computing, 17(5):967-974, 1988.
[17] Rolfdieter Frank and Harald Riede. Hyperplane sections of the n-dimensional cube. The American Mathematical Monthly, 119(10):868-872, 2012.
[18] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2021.
[19] Harold W Kuhn. On games of fair division, essays in mathematical economics in honor of oskar morgenstern, edited by martin shubik, 1967.
[20] Maria Kyropoulou, Josué Ortega, and Erel Segal-Halevi. Fair cake-cutting in practice. In Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019, pages 547-548, 2019.
[21] László Lovász and Santosh S. Vempala. Hit-and-run from a corner. SIAM J. Comput., 35(4):985-1005, 2006.
[22] M Meyer and S Reisner. Characterization of affinely-rotation-invariant log-concave measures by section-centroid location. In Geometric Aspects of Functional Analysis: Israel Seminar (GAFA) 1989-90, pages 145-152. Springer, 1991.
[23] Antonio Nicolò and Yan Yu. Strategic divide and choose. Games and Economic Behavior, 64(1):268-289, 2008.
[24] Brian R Powers. An analysis of dual-issue final-offer arbitration. International Journal of Game Theory, 48(1):81-108, 2019.
[25] Ariel D. Procaccia. Cake cutting: not just child's play. Communications of the ACM, 56(7):7887, 2013.
[26] Milton Weinstein and Richard Zeckhauser. Critical ratios and efficient allocation. Journal of Public Economics, 2(2):147-157, 1973.
[27] H. Peyton Young. Equity - in theory and practice. Princeton University Press, 1995.

## A Proof of Lemma 3.1

To prove (i), observe that the average utility from the chooser's two options is

$$
\frac{1}{2}\left(\sum_{i=1}^{n} p_{i} g_{i}^{C}+\sum_{i=1}^{n}\left(1-p_{i}\right) g_{i}^{C}\right)=\frac{1}{2} \sum_{i=1}^{n} g_{i}^{C},
$$

so at least one of the options must yield at least this utility. To prove (ii), observe that the division $p=\left(\frac{1}{2}, \frac{1}{2}, \ldots, \frac{1}{2}\right)$ yields utility $\sum_{i=1}^{n}\left(\frac{1}{2}\right) g_{i}^{D}=\frac{1}{2} \sum_{i=1}^{n} g_{i}^{D}$ when the chooser picks pile 2 and utility $\sum_{i=1}^{n}\left(1-\frac{1}{2}\right) g_{i}^{D}=\frac{1}{2} \sum_{i=1}^{n} g_{i}^{D}$ when the chooser picks pile 1 . As the divider always has a strategy guaranteeing utility $\frac{1}{2} \sum_{i=1}^{n} g_{i}^{D}$, at equilibrium he must receive at least that utility in expectation.

## B Proof of Lemma 3.2

We first show that optimal divisions exist. Using Equation (2), We may express the divider's expected utility as

$$
\begin{align*}
\mathbb{E}\left[u^{D}\right] & =P\left(\sum_{i=1}^{n}\left(1-p_{i}\right) g_{i}^{D}\right)+(1-P)\left(\sum_{i=1}^{n} p_{i} g_{i}^{D}\right) \\
& =P\left(\sum_{i=1}^{n}\left(1-\frac{q_{i}}{2}-\frac{1}{2}\right) g_{i}^{D}\right)+(1-P)\left(\sum_{i=1}^{n}\left(\frac{q_{i}}{2}+\frac{1}{2}\right) g_{i}^{D}\right) \\
& =P \sum_{i=1}^{n} \frac{g_{i}^{D}}{2}-P \sum_{i=1}^{n} \frac{g_{i}^{D} q_{i}}{2}+\sum_{i=1}^{n} \frac{g_{i}^{D} q_{i}}{2}+\sum_{i=1}^{n} \frac{g_{i}^{D}}{2}-P \sum_{i=1}^{n} \frac{g_{i}^{D} q_{i}}{2}-P \sum_{i=1}^{n} \frac{g_{i}^{D}}{2} \\
& =\sum_{i=1}^{n} \frac{g_{i}^{D}}{2}+\left(\frac{1}{2}-P\right) \sum_{i=1}^{n} q_{i} g_{i}^{D} . \tag{10}
\end{align*}
$$

The optimal utility is attained by maximizing Equation (10) over the variables $q_{1}, q_{2}, \ldots, q_{i} \in[-1,1]$. Let $u^{*}$ denote the supremum of this optimal utility, and consider a sequence of divisions $q^{1}, q^{2}, q^{3}, \ldots$ whose expected utilities converge to $u^{*}$. Let $q^{*} \in[-1,1]$ be the division in the limit of a convergent subsequence. We claim that $q^{*}$ yields the optimal utility $u^{*}$. Indeed, since Equation (10) is a continuous function of $q$ and $P$, this could only fail if there were a discontinuity in the function $P(q)$ at $q^{*}$. Such a discontinuity must be the result of a set of chooser types $S$ of positive probability mass where all types in $S$ pick the divider's strictly preferred pile at $q^{*}$ and the other pile after an arbitrarily small deviation from $q^{*}$. But this means that all chooser types in $S$ are indifferent between the two piles, which contradicts our assumption that an indifferent chooser breaks her indifference in favor of the divider. Thus, $q^{*}$ attains the optimal utility $u^{*}$.

If a given division does not satisfy Equation (3), then sending each $q_{i} \mapsto-q_{i}$ will satisfy it. This corresponds to sending $p_{i} \mapsto 1-p_{i}$, so it is an equivalent division up to renaming the piles. Thus, it is without loss of generality to assume (3) holds in an optimal division.

If $P>\frac{1}{2}$, then, since the final sum in Equation (10) is nonnegative by (3), we have

$$
\mathbb{E}\left[u^{D}\right] \leq \sum_{i=1}^{n} \frac{g_{i}^{D}}{2},
$$

so the divider is no better off than his proportionality guarantee. Therefore, the divider is at least equally well-off setting $p=\left(\frac{1}{2}, \frac{1}{2}, \ldots, \frac{1}{2}\right)$, in which case it is without loss of generality that $P \leq \frac{1}{2}$.

To prove the final statement, simply observe that the final term in Equation (10) is nonzero if and only if the inequalities in both Equations (3) and (4) are strict.

## C Proof of Lemma 3.5

Observe that each division $p_{1}, p_{2}, \ldots, p_{n}$ computed by Algorithm 1 is valid, since $0 \leq p_{i} \leq 1$ if and only if $-1 \leq q_{i} \leq 1$, which is enforced by the first constraint of $\mathcal{C}_{P}$. We claim that, on each iteration of the main loop, the division $p_{1}, p_{2}, \ldots, p_{n}$ computed by Algorithm 1 achieves an interim expected utility of $\mathbb{E}\left[u^{D}\right] \geq u_{P}$. (In fact, the utility will be exactly $u_{P}$, but equality is not necessary to prove.)

The chooser weakly prefers pile 1 if and only if

$$
\sum_{i=1}^{n} g_{i}^{C} p_{i} \geq \sum_{i=1}^{n} g_{i}^{C}\left(1-p_{i}\right)
$$

Note that this is equivalent to

$$
s:=\sum_{i=1}^{n} g_{i}^{C} q_{i} \geq 0
$$

Since each $g_{i}^{C}$ follows a normal distribution with mean $\mu_{i}$ and variance $\sigma_{i}^{2}$, we know that $s$ follows a normal distribution with mean

$$
\sum_{i=1}^{n} \mu_{i} q_{i}
$$

and variance

$$
\sum_{i=1}^{n} \sigma_{i}^{2} q_{i}^{2}
$$

Hence, the probability that $s \geq 0$ is given by

$$
1-\Phi\left(\frac{0-\sum_{i=1}^{n} \mu_{i} q_{i}}{\sqrt{\sum_{i=1}^{n} \sigma_{i}^{2} q_{i}^{2}}}\right)=\Phi\left(\frac{\sum_{i=1}^{n} \mu_{i} q_{i}}{\sqrt{\sum_{i=1}^{n} \sigma_{i}^{2} q_{i}^{2}}}\right) .
$$

Since the algorithm computes optimal $q_{1}, q_{2}, \ldots, q_{n}$ on each iteration of the main loop to satisfy the
third constraint, we know that this probability is at most $P$. Therefore,

$$
\begin{aligned}
\mathbb{E}\left[u^{D}\right] & =\operatorname{Pr}[\text { chooser picks pile } 1] \sum_{i=1}^{n} g_{i}^{D}\left(1-p_{i}\right)+\operatorname{Pr}[\text { chooser picks pile } 2] \sum_{i=1}^{n} g_{i}^{D} p_{i} \\
& =\sum_{i=1}^{n} g_{i}^{D} p_{i}+\operatorname{Pr}\left[\text { chooser picks pile 1] }\left(\sum_{i=1}^{n} g_{i}^{D}\left(1-p_{i}\right)-\sum_{i=1}^{n} g_{i}^{D} p_{i}\right)\right. \\
& \geq \sum_{i=1}^{n} g_{i}^{D} p_{i}+P\left(\sum_{i=1}^{n} g_{i}^{D}\left(1-p_{i}\right)-\sum_{i=1}^{n} g_{i}^{D} p_{i}\right)
\end{aligned}
$$

(since the second constraint ensures the term in parentheses is nonpositive)
$=\sum_{i=1}^{n} g_{i}^{D}\left(P\left(1-p_{i}\right)+(1-P) p_{i}\right)$
$=\sum_{i=1}^{n} g_{i}^{D}\left(P\left(1-\left(\frac{q_{i}}{2}+\frac{1}{2}\right)\right)+(1-P)\left(\frac{q_{i}}{2}+\frac{1}{2}\right)\right)$
$=\sum_{i=1}^{n} \frac{g_{i}^{D}}{2}\left(P\left(1-q_{i}\right)+(1-P)\left(1+q_{i}\right)\right)$
$=u_{P}$.
Thus, the claim is proved.
Let $\left(p_{1}^{*}, p_{2}^{*}, \ldots, p_{n}^{*}\right)$ denote an optimal division from Lemma 3.2, yielding interim expected utility $u^{*}$, and let $\left(q_{1}^{*}, q_{2}^{*}, \ldots, q_{n}^{*}\right)$ denote the respective auxiliary variables for this division (i.e., obtained from Equation (1)). Recall that, in this optimal division, the divider weakly prefers pile 1, and the probability that the chooser picks pile 1 is $P^{*} \leq \frac{1}{2}$. Therefore, on some iteration of Algorithm 1,

$$
P-\delta \leq P^{*} \leq P
$$

Let $\left(q_{1}, q_{2}, \ldots, q_{n}\right)$ denote the optimal solution to $\mathcal{C}_{P}$ on this iteration, with optimal value $u_{P}$, and let $\left(p_{1}, p_{2}, \ldots, p_{n}\right)$ denote the corresponding division.

Observe that $\left(q_{1}^{*}, q_{2}^{*}, \ldots, q_{n}^{*}\right)$ is feasible for $\mathcal{C}_{P}$. To see this, note that the first constraint is satisfied by the fact that it corresponds to a valid division with each $p_{i} \in[0,1]$. The second constraint is satisfied because we are assuming the divider prefers pile 1. Finally, for the third constraint, since $P^{*} \leq P$ implies $\Phi\left(P^{*}\right) \leq \Phi(P)$, we have

$$
\frac{\sum_{i=1}^{n} \mu_{i} q_{i}^{*}}{\sqrt{\sum_{i=1}^{n} \sigma_{i}^{2}\left(q_{i}^{*}\right)^{2}}} \leq \Phi\left(P^{*}\right) \leq \Phi(P) .
$$

Thus, $\left(q_{1}^{*}, q_{2}^{*}, \ldots, q_{n}^{*}\right)$ is a feasible solution for $\mathcal{C}_{P}$.
Therefore, denoting the objective function of $\mathcal{C}_{P}$ by $f_{P}$, we have that the utility of the optimal
solution returned by the algorithm is

$$
\begin{aligned}
& \mathbb{E}\left[u^{D}\right] \geq u_{P} \quad \text { (from the previous claim) } \\
& =f_{P}\left(q_{1}, q_{2}, \ldots, q_{n}\right) \\
& \geq f_{P}\left(q_{1}^{*}, q_{2}^{*}, \ldots, q_{n}^{*}\right) \quad\left(\operatorname{since}\left(q_{1}^{*}, q_{2}^{*}, \ldots, q_{n}^{*}\right) \text { is feasible and }\left(q_{1}, q_{2}, \ldots, q_{n}\right)\right. \text { is optimal) } \\
& =\sum_{i=1}^{n} \frac{g_{i}^{D}}{2}\left(P\left(1-q_{i}^{*}\right)+(1-P)\left(1+q_{i}^{*}\right)\right) \\
& =\sum_{i=1}^{n} \frac{g_{i}^{D}}{2}\left(1+q_{i}^{*}\right)-P \sum_{i=1}^{n} g_{i}^{D} q_{i}^{*} \\
& =\sum_{i=1}^{n} \frac{g_{i}^{D}}{2}\left(1+q_{i}^{*}\right)-P^{*} \sum_{i=1}^{n} g_{i}^{D} q_{i}^{*}-\left(P-P^{*}\right) \sum_{i=1}^{n} g_{i}^{D} q_{i}^{*} \\
& =\sum_{i=1}^{n} \frac{g_{i}^{D}}{2}\left(1+2 p_{i}^{*}-1\right)-P^{*} \sum_{i=1}^{n} g_{i}^{D}\left(2 p_{i}^{*}-1\right)-\left(P-P^{*}\right) \sum_{i=1}^{n} g_{i}^{D} q_{i}^{*} \\
& =P^{*} \sum_{i=1}^{n} g_{i}^{D}\left(1-p_{i}^{*}\right)+\left(1-P^{*}\right) \sum_{i=1}^{n} g_{i}^{D} p_{i}^{*}-\left(P-P^{*}\right) \sum_{i=1}^{n} g_{i}^{D} q_{i}^{*} \\
& =\operatorname{Pr}\left[\text { chooser picks 1] } \sum_{i=1}^{n} g_{i}^{D}\left(1-p_{i}^{*}\right)+\operatorname{Pr}[\text { chooser picks } 2] \sum_{i=1}^{n} g_{i}^{D} p_{i}^{*}-\left(P-P^{*}\right) \sum_{i=1}^{n} g_{i}^{D} q_{i}^{*}\right. \\
& =u^{*}-\left(P-P^{*}\right) \sum_{i=1}^{n} g_{i}^{D} q_{i}^{*} \\
& \geq u^{*}-\left|\left(P-P^{*}\right)\right| \sum_{i=1}^{n}\left|g_{i}^{D}\right|\left|q_{i}^{*}\right| \\
& \geq u^{*}-\delta \sum_{i=1}^{n}\left|g_{i}^{D}\right| \quad\left(\text { since } P-\delta \leq P^{*} \text { and each }\left|q_{i}\right| \leq 1\right) \\
& =u^{*}-\gamma
\end{aligned}
$$

as desired.

## D Proof of Theorem 3.10

For any $0 \leq P \leq \frac{1}{2}$ and $V>0$, let

$$
S_{P, V}:=\left\{q \in[-1,1]^{n} \mid \operatorname{Pr}_{g^{C} \sim \overline{\mathcal{G}}^{C}}\left[q \cdot g^{C} \geq 0\right] \leq P \text { and } q \cdot g^{D} \geq V\right\}
$$

This is the set of divisions for which the chooser picks Pile 1 with probability at most $P$ and Pile 1 is worth at least $V$ more to the divider than Pile 2 . We will describe an algorithm, running in polynomial time in the size of the input, $\frac{1}{\varepsilon}$, and $\frac{1}{\delta}$, that either finds a division $q \in S_{P+\varepsilon, V}$ or reports that $S_{P, V}$ is empty, and is correct with probability $\geq 1-\delta$. It is straightforward to see that such an algorithm can be called repeatedly to give the desired FPRAS: we apply the same search over $P$ as in Algorithm 1. However, instead of solving the quadratic program, we repeatedly test feasibility of $S_{P, V}$, using a binary-search to find the maximal $V$.

Consider the following randomized separation oracle $\mathcal{O}$ for $S_{P, V}$. Given a queried point $q$, first we test whether $q \cdot g^{D} \geq V$. If not, then we return this inequality as a violated constraint.

Otherwise, we empirically estimate the probability $r$ that the chooser picks Pile 1 by repeatedly sampling many chooser values $g^{C}$ and counting the number of times $q \cdot g^{C} \geq 0$. If the empirical estimate $\widehat{r}$ is at most $P+\frac{2 \varepsilon}{3}$, then we return true. Otherwise, we compute an estimate $\widehat{m}$ of the centroid $m$ of the set

$$
M:=\left\{g^{C} \in \prod_{i=1}^{n}\left[a_{i}, b_{i}\right] \mid q \cdot g^{C}=0\right\}
$$

and return the violated linear constraint that $q \cdot \widehat{m}<0$. Computing the approximate centroid $\widehat{m}$ can be accomplished by repeatedly sampling many uniformly random points from $M$ and taking the average, which can be done in polynomial time using the hit-and-run random walk [21]. We omit the tedious details concerning the sample complexity of estimating $r$ and $m$. All we require is that, for any $\delta^{\prime}>0$, we can, in polynomial time in $\frac{1}{\varepsilon \delta^{\prime}}$, ensure with probability at least $1-\delta^{\prime}$ that $\widehat{r}$ is within $\frac{\varepsilon}{3}$ of the true probability $r$ and $\widehat{m}$ is within

$$
\frac{\varepsilon}{3} \min _{1 \leq i \leq n}\left(b_{i}-a_{i}\right)
$$

of the true centroid $m$ (under the $L_{1}$ norm). We say $\mathcal{O}$ misses if one of these two conditions is false.

We run the ellipsoid algorithm using separation oracle $\mathcal{O}$ until it returns true. This takes a number of iterations $s$ that is at most a polynomial function of the size of the input and $\frac{1}{\varepsilon}$. Letting $\delta^{\prime}=\frac{\delta}{s}$, each call to $\mathcal{O}$ will take polynomial time, and we know by the union bound that the probability some iteration misses is at most $\delta$. Thus, to show that the algorithm is correct, we just need to show that $\mathcal{O}$ returns a correct result whenever it does not miss. There are three cases to consider, depending on what $\mathcal{O}$ returns when queried on division $q$.

If $\mathcal{O}(q)$ returns true, then $q \cdot g^{D} \leq V$ and $\widehat{r} \leq P+\frac{2 \varepsilon}{3}$. Assuming $\mathcal{O}$ does not miss, we will have $|\widehat{r}-r| \leq \frac{\varepsilon}{3}$, in which case $\widehat{r} \leq P+\frac{2 \varepsilon}{3}+\frac{\varepsilon}{3}=P+\varepsilon$ by the triangle inequality. Hence, the algorithm has found a division $q \in S_{P+\varepsilon, V}$ as desired.

If $\mathcal{O}(q)$ returns that $q$ violates the inequality $q \cdot g^{D} \geq V$, then it is always correct (in fact, there was no way for it to miss). It has determined that $q$ does not satisfy this constraint, and obviously any $q^{\prime} \in S_{P, V}$ does satisfy it.

Finally, consider the case where $\mathcal{O}(q)$ returns that $q$ violates the inequality $q \cdot \widehat{m}<0$. Observe that $q$ does indeed violate this inequality because $\widehat{m} \in M$. All that remains to show is that all $q^{\prime} \in S_{P, V}$ do satisfy it. We show the contrapositive, that an arbitrary $q^{\prime}$ violating the inequality is not in $S_{P, V}$. So suppose $q^{\prime} \cdot \widehat{m} \geq 0$. Without loss of generality, assume that $q_{n}$ has the largest magnitude out of all $q_{i}$ for all $1 \leq i \leq n$. Note that $q_{n} \neq 0$, for otherwise $q$ is the all-zero vector, so $q \cdot g^{D} \geq V$ would have been violated, as we are assuming $V>0$. Since $q^{\prime} \cdot \widehat{m}$ is positive, we have
that

$$
\begin{aligned}
\operatorname{Pr}_{g^{C} \sim \overline{\mathcal{G}}^{C}}\left[q^{\prime} \cdot g^{C} \geq 0\right] & \geq \operatorname{Pr}_{g^{C} \sim \overline{\mathcal{G}}^{C}}\left[q^{\prime} \cdot g^{C}-q^{\prime} \cdot \widehat{m} \geq 0\right] \\
& =\operatorname{Pr}_{g^{C} \sim \overline{\mathcal{G}}^{C}}\left[q^{\prime} \cdot g^{C} \geq q^{\prime} \cdot \widehat{m}\right] \\
& \geq \operatorname{Pr}_{g^{C} \sim \overline{\mathcal{G}}^{C}}\left[q^{\prime} \cdot g^{C} \geq q^{\prime} \cdot m\right]-\operatorname{Pr}_{g^{C} \sim \overline{\mathcal{G}}^{C}}\left[q^{\prime} \cdot \widehat{m} \geq q^{\prime} \cdot g^{C} \geq q^{\prime} \cdot m\right] \\
& =\operatorname{Pr}_{g^{C} \sim \overline{\mathcal{G}}^{C}}\left[q^{\prime} \cdot\left(g^{C}-m\right) \geq 0\right]-\operatorname{Pr}_{g^{C} \sim \overline{\mathcal{G}}^{C}}\left[g_{n}^{C} \in J\right] \text { where } \\
& J:=\left[\frac{1}{q_{n}^{\prime}}\left(\sum_{i=1}^{n} q_{i}^{\prime} \widehat{m}_{i}-\sum_{i=1}^{n-1} q_{i}^{\prime} g_{i}^{C}\right), \frac{1}{q_{n}^{\prime}}\left(\sum_{i=1}^{n} q_{i}^{\prime} m_{i}-\sum_{i=1}^{n-1} q_{i}^{\prime} g_{i}^{C}\right)\right] \\
& >\left(P+\frac{\varepsilon}{3}\right)-\frac{\varepsilon}{3} \quad \quad \text { (see justification below) } \\
& =P .
\end{aligned}
$$

Therefore, $q^{\prime} \notin S_{P, V}$. It remains to justify the final inequality above, which follows from the following two inequalities:

$$
\begin{gather*}
\operatorname{Pr}_{g^{C} \sim \overline{\mathcal{G}}^{C}}\left[q^{\prime} \cdot\left(g^{C}-m\right) \geq 0\right]>P+\frac{\varepsilon}{3}  \tag{11}\\
\operatorname{Pr}_{g^{C} \sim \overline{\mathcal{G}}^{C}}\left[g_{n}^{C} \in J\right] \leq \frac{\varepsilon}{3} \tag{12}
\end{gather*}
$$

To prove Inequality (11), let $H_{1}$ and $H_{2}$ be the ( $n-1$ )-dimensional planes

$$
H_{1}:=\left\{g^{C} \in \mathbb{R}^{n} \mid q^{\prime} \cdot\left(g^{C}-m\right)=0\right\}, \quad H_{2}:=\left\{g^{C} \in \mathbb{R}^{n} \mid q \cdot g^{C}=0\right\}
$$

with positive sides defined as

$$
H_{1}^{+}:=\left\{g^{C} \in \mathbb{R}^{n} \mid q^{\prime} \cdot\left(g^{C}-m\right) \geq 0\right\}, \quad H_{2}^{+}:=\left\{g^{C} \in \mathbb{R}^{n} \mid q \cdot g^{C} \geq 0\right\}
$$

Recall that $r$ is the true probability that $q \cdot g^{C} \geq 0$, and $\widehat{r}$ is the estimate of $r$. Applying Theorem 3.8 part (ii) to $H_{1}$, we know that $m \in K_{r}$. Since $q^{\prime} \cdot(m-m)=0$, we have $m \in H_{2}^{+}$as well, so $m \in H_{2}^{+} \cap K_{r}$. Hence, applying Theorem 3.8 part (i) (backward direction) to $H_{2}$, we derive that

$$
\operatorname{Pr}_{g^{C} \sim \overline{\mathcal{G}}^{C}}\left[q^{\prime} \cdot\left(g^{C}-m\right) \geq 0\right]=\operatorname{Pr}_{g^{C} \sim \overline{\mathcal{G}}^{C}}\left[g^{C} \in H_{2}^{+}\right] \geq r .
$$

Assuming $\mathcal{O}$ did not miss, we additionally know that $|r-\widehat{r}| \leq \frac{\varepsilon}{3}$ and $\widehat{r}>P+\frac{2 \varepsilon}{3}$, so Inequality (11) follows from the triangle inequality.

Finally, we prove Inequality (12) holds under any fixed values of $g_{i}^{C}$ for $1 \leq i \leq n-1$, letting $g_{n}^{C}$ be the only random variable. Using the assumptions that $q_{n}^{\prime}$ has the largest magnitude and that $\mathcal{O}$ did not miss, we know that

$$
|J|=\sum_{i=1}^{n} \frac{q_{i}^{\prime}}{q_{n}^{\prime}}\left(\widehat{m}_{i}-m_{i}\right) \leq \sum_{i=1}^{n}\left|\widehat{m}_{i}-m_{i}\right| \leq \frac{\varepsilon}{3}\left(b_{n}-a_{n}\right) .
$$

Therefore, since $g_{n}^{C} \sim\left[a_{n}, b_{n}\right]$, the probability $g_{i}^{C} \in J$ is at most $\frac{\varepsilon}{3}$, so Inequality (12) holds.
This proves that the separation oracle is correct with probability $1-\delta$, and thus so is the algorithm.

## E Proof of Lemma 3.5

Suppose $\overline{\mathcal{G}}^{C}$ is an arbitrary distribution supported on a finite set $\left\{x_{1}, x_{2}, \ldots, x_{\ell}\right\} \subseteq \mathbb{R}^{n}$. For each $1 \leq j \leq \ell$, we write $r_{j}$ for the probability that $g^{C}=x_{j}$. The main idea for finding an optimal division is similar to that of Algorithm 1: we try to guess the set of chooser types who pick pile 1 ; we then use linear programming to compute the optimal division with respect to the additional constraints that entails. (Since there are an exponential number of subsets in contention, this means the algorithm is only practical when the number of types $\ell$ is small.) Specifically, consider the following algorithm.

```
Algorithm 2: Computes an optimal division given the divider's values and an arbitrary
discrete prior for the chooser's values.
    Input: Divider values \(g_{1}^{D}, g_{2}^{D}, \ldots, g_{n}^{D}\) (not all zero), possible chooser value vectors
            \(x_{1}, x_{2}, \ldots, x_{\ell} \in \mathbb{R}^{n}\), and corresponding probabilities \(r_{1}, r_{2}, \ldots, r_{\ell}\).
    Output: An optimal division \(p_{1}, p_{2}, \ldots, p_{n}\).
    \(u \leftarrow-\infty\)
    for \(S \subseteq[\ell]\) do
        \(P \leftarrow \sum_{j \in S} r_{j}\)
        if \(P \leq \frac{1}{2}\) then
            \(u^{\prime}, q_{1}, q_{2}, \ldots, q_{i} \leftarrow\) optimal solution to the following linear program \(\mathcal{C}_{S}\) :
                    \(\operatorname{maximize} \quad u^{\prime}=\sum_{i=1}^{n} \frac{g_{i}^{D}}{2}\left(P\left(1-q_{i}\right)+(1-P)\left(1+q_{i}\right)\right)\)
                \(\begin{array}{lll}\text { subject to } & -1 \leq q_{i} \leq 1 & \text { for all } 1 \leq i \leq n, \\ & \sum_{i=1}^{n}\left(x_{j}\right)_{i} q_{i} \leq 0 & \text { for all } j \in[\ell] \backslash S\end{array}\)
            if \(u^{\prime}>u\) then
            \(u \leftarrow u^{\prime}\)
                for \(i \leftarrow 1,2, \ldots, n\) do
                \(p_{i} \leftarrow \frac{q_{i}+1}{2}\)
            end
            end
        end
    end
    return \(\left(p_{1}, p_{2}, \ldots, p_{n}\right)\)
```

Let $p^{*}=\left(p_{1}^{*}, p_{2}^{*}, \ldots, p_{n}^{*}\right)$ denote an optimal division from Lemma 3.2 , with divider utility $u^{*}$ and probability $P^{*}$ that the chooser picks pile 1 . Let $S^{*} \subseteq[\ell]$ denote the set of chooser types $j$ such that, when the chooser has value vector $g^{C}=x_{j}$, she chooses pile 1 given the division $p^{*}$. Note that the probability the chooser picks pile 1 given $p^{*}$ is $P^{*}=\sum_{j \in S^{*}} r_{j} \leq \frac{1}{2}$. It is not too hard to see that $\mathcal{C}_{S}$ finds the optimal division given the additional constraint that all types $j \notin S$ pick pile 2 , assuming further that all types $j \in S$ will pick pile 1 (which can only lower the objective value). In particular, this means that $u^{\prime}$ is always at least the utility from some feasible division, and on the iteration of the main loop when $S=S^{*}$, Algorithm 2 will successfully find $p=p^{*}$, and the optimal utility will be $u=u^{*}$. The claim about running time follows from the observation that each $\mathcal{C}_{S}$ is a linear program.

## F Proof of Lemma 4.2

Let $q$ be an optimal division from Lemma 3.2. If $q$ achieves a higher-than-proportional utility, then there must be some good $i$ such that $\left|q_{i}\right|>0$. Let $i^{*}$ be a good maximizing $\left|q_{i^{*}}\right|$. If $\left|q_{i^{*}}\right|=1$, then good $i^{*}$ is put entirely into one of the two piles; otherwise, we claim that the divider can improve the division by linearly scaling $q$, dividing each component by $\left|q_{i^{*}}\right|$. Since

$$
P=\operatorname{Pr}\left[\sum_{i=1}^{n} q_{i} g_{i}^{C}>0\right]=\operatorname{Pr}\left[\sum_{i=1}^{n} \frac{q_{i}}{\left|q_{i^{*}}\right|} g_{i}^{C}>0\right],
$$

the chooser is just as likely to pick pile 1 , but now the divider obtains a greater utility from pile 1 , so he receives a higher expected utility.

## G Tightness of Theorem 5.1

In this section we provide counterexamples to possible ways that one might hope to strengthen/generalize Theorem 5.1.

- The condition on $\mathcal{D}$ that the critical ratios are not always identical. Otherwise the divider may not be able to beat his proportionality guarantee. For instance, in the extreme case where the values of the two goods are always identical, both roles confer the same utility, which would violate both parts of the theorem.
- The assumption that values are not correlated for $n=2$. Suppose $\mathcal{D}$ is a distribution where the divider faces no uncertainty of the chooser's value after seeing his own value. Then the divider can extract all of the surplus from having different values, leaving the chooser with utility equal to her proportionality guarantee. Hence, correlations can make the divider better off, even for only two goods.
- Positivity of $\mathcal{D}$. If $\mathcal{D}$ is the uniform distribution on $[-1,1]^{2}$, one can see that $P=\frac{1}{2}$ in any division, so the divider cannot surpass his proportionality guarantee utility for any number of goods. However, for large $n$, the chooser will certainly reap more than her proportional utility (assuming the divider breaks his indifference between divisions in a way that benefits the chooser).
- The condition in Equation (9). Roughly, this says that the values from $\mathcal{D}$ are typically not extremely small. To see why this assumption is necessary, suppose the divider's value for each good is 0 with probability $90 \%$ and 1 with probability $10 \%$, while the chooser's value is uniformly distributed on $[0,1]$. For a large number of goods, the divider will place the roughly $10 \%$ of the goods for which he has high-value in pile 1 and the remaining roughly $90 \%$ of the goods in pile 2. By choosing pile 2, the chooser obtains an ex-ante expected utility of $0.9 \cdot 0.5=0.45$ per good. If she were the divider, she would have to create piles of roughly even size (with the preferred pile 1 containing the best half of the goods, slightly smaller than pile 2 ), only receiving an ex-ante utility of $0.5 \cdot 0.75=0.375$ per good. Thus, in this example, it is favorable for the first player to be the chooser. However, this is only because the second player sometimes draws the value zero. As we show in the proof, as long as the divider's value is not zero, even the slightest bit greater than zero, the optimal division leaves the chooser with only her proportionality utility guarantee.
- The chooser is better off with two goods but not always with three goods. Suppose each of the three goods have value 1 with $99 \%$ probability and and 300 with $1 \%$ probability. There are three cases to consider.
- With probability $\left(\frac{99}{100}\right)^{6} \approx 0.941$, all six valuations will be the low value of 1 . The divider will divide the goods evenly and the chooser will pick an arbitrary pile, with both players gaining the same utility.
- With probability $6\left(\frac{99}{100}\right)^{5}\left(\frac{1}{100}\right) \approx 0.057$, there will be exactly one high-valued good. If it is the divider's value, the divider will optimally place it alone, entirely in pile 1 and most of the other two goods in pile 2 . The divider's expected utility is at least $\left(\frac{99}{100}\right) \cdot 300=297$ and the chooser's expected utility is at most 2 . If the chooser has the high valuation, the divider will evenly divide the goods between the piles, so there is a chance the chooser's favorite good will be fractionally allocated. Specifically, some good must be divided with at least $\frac{1}{4}$ in each pile. So the chooser's expected utility from the high-value good is at most $\left(\frac{2}{3}\right) \cdot 300+\left(\frac{1}{3}\right) \cdot\left(\frac{3}{4}\right) \cdot 300=275$, for a total utility of at most 278 . Thus, in this case, the expected utility of the divider is higher than that of the chooser by at least

$$
\frac{1}{2}(297-2)+\frac{1}{2}(0-278)=8.5
$$

- With probability $15\left(\frac{99}{100}\right)^{4}\left(\frac{1}{100}\right)^{2} \approx 0.00144$, there will be exactly two high-valued goods. We can upper bound the chooser's utility in this case by the expected value of all goods to the chooser, which is $300+1+1=302$.
- With the remaining probability of $1-\left(\frac{99}{100}\right)^{6}-6\left(\frac{99}{100}\right)^{5}\left(\frac{1}{100}\right)-15\left(\frac{99}{100}\right)^{4}\left(\frac{1}{100}\right)^{2} \approx 0.00002$, more than two goods have high value, in which case we may upper bound the chooser's utility by the maximal value of the sum of all goods, which is 900 .

In total, the expected utility of the divider minus that of the chooser is at least

$$
0.057 \cdot 8.5-0.00144 \cdot 302-0.00002 \cdot 900 \approx 0.03 \geq 0
$$

## H Proof of Theorem 5.1 part (ii)

Assume without loss of generality that the support of $\mathcal{D}$ is contained in $[0,1] \times[0,1]$. Note that this means the maximum possible utility for either player is at most $n$. Let

$$
B:=\frac{1}{2} \underset{\left(g^{1}, g^{2}\right) \sim \mathcal{D}}{\mathbb{E}}\left[g^{1}\right]
$$

We will show that, for all sufficiently large $n$, the player whose values are drawn according to the first component of $\mathcal{D}$ attains an expected utility of at least $n(B+\varepsilon)$ as the divider for some constant $\varepsilon>0$, but attains utility strictly less than this as the chooser for any $\varepsilon>0$.

We begin with the divider's utility. Since the critical ratio of a good is not always the same value, there must be a pair of values $b_{1}>b_{2}>0$, such that the critical ratio of a good is at least $b_{1}$ with some nonzero probability, and at most $b_{2}$ with some nonzero probability. By similar reasoning, we can write $b_{1}=\frac{c_{1}}{c_{3}}$ such that, with nonzero probability, the numerator of the critical ratio is bounded below by $c_{1}$ and the denominator is bounded above by $c_{3}$; likewise, we can write $b_{2}=\frac{c_{2}}{c_{4}}$ such that, with nonzero probability, the numerator of the critical ratio is bounded above by $c_{2}$ and the denominator is bounded below by $c_{4}$. Note that $b_{1}>b_{2}$ implies

$$
\begin{equation*}
c_{1} c_{4}>c_{2} c_{3} \tag{13}
\end{equation*}
$$

By independence of the values of goods $i$ and $i+1$, it follows that, with some probability $c_{5}>0$, good $i$ will have a higher critical ratio than good $i+1$ as according to these bounds:

$$
\begin{equation*}
\operatorname{Pr}_{\substack{g_{i}^{D}, g_{c}^{C} \sim \mathcal{D} \\ g_{i+1}^{D}, g_{i+1}^{C} \sim \mathcal{D}}}\left[g_{i}^{D} \geq c_{1}, g_{i+1}^{D} \leq c_{2}, \mathbb{E}\left[g_{i}^{C}\right] \leq c_{3}, \text { and } \mathbb{E}\left[g_{i+1}^{C}\right] \geq c_{4}\right] \geq c_{5} \tag{14}
\end{equation*}
$$

We let $S$ be the set of odd $i<n$ for which the condition in Equation 14 is satisfied. Consider the following division, which is reminiscent of the proof of Proposition 3.3:

$$
q_{i}:=\left\{\begin{array}{ll}
\frac{c_{2}+c_{4}}{2 \alpha} & \text { if } i \in S \\
-\frac{c_{1}+c_{3}}{2 \alpha} & \text { if } i-1 \in S \\
0 & \text { otherwise }
\end{array} \quad \text { where } \quad \alpha:=\max \left\{\frac{c_{1}+c_{3}}{2}, \frac{c_{2}+c_{4}}{2}\right\}\right.
$$

The divider strictly prefers pile 1 by

$$
\begin{aligned}
\sum_{i=1}^{n} q_{i} g_{i}^{D} & =\sum_{i \in S}\left(\frac{c_{2}+c_{4}}{2 \alpha} g_{i}^{D}-\frac{c_{1}+c_{3}}{2 \alpha} g_{i+1}^{D}\right) \\
& \geq \sum_{i \in S}\left(\frac{1}{2 \alpha}\left(c_{2}+c_{4}\right) c_{1}-\frac{1}{2 \alpha}\left(c_{1}+c_{3}\right) c_{2}\right) \\
& =\frac{|S|}{2 \alpha}\left(c_{1} c_{4}-c_{2} c_{3}\right)
\end{aligned}
$$

which is a constant greater than zero. Likewise, the interim expectation of the difference between the two piles' value for the chooser is

$$
\begin{aligned}
\mathbb{E}\left[\sum_{i=1}^{n} q_{i} g_{i}^{C}\right] & =\sum_{i \in S}\left(\frac{c_{2}+c_{4}}{2 \alpha} \mathbb{E}\left[g_{i}^{C}\right]-\frac{c_{1}+c_{3}}{2 \alpha} \mathbb{E}\left[g_{i+1}^{C}\right]\right) \\
& \leq \sum_{i \in S}\left(\frac{c_{2}+c_{4}}{2 \alpha} c_{3}-\frac{c_{1}+c_{3}}{2 \alpha} c_{4}\right) \\
& =\frac{|S|}{2 \alpha}\left(c_{2} c_{3}-c_{1} c_{4}\right),
\end{aligned}
$$

which is a constant less than zero. Thus, by standard concentration bounds, for large enough $n$ the probability the chooser picks pile 1 converges to zero, so the divider's utility gains beyond the proportionality guarantee converges to

$$
\frac{\mathbb{E}[|S|]}{4 \alpha}\left(c_{1} c_{4}-c_{2} c_{3}\right)=\frac{n c_{5}\left(c_{1} c_{4}-c_{2} c_{3}\right)}{4 \alpha} .
$$

Thus, letting $\varepsilon:=\frac{c_{5}\left(c_{1} c_{4}-c_{2} c_{3}\right)}{5 \alpha}>0$, the desired conclusion follows.
We now compute the utility of the first component player if they are the chooser. We need to be slightly more formal in stating the concentration bounds required for this case. Specifically, given $\varepsilon>0$, we choose $n$ large enough so that $n \geq \frac{4}{\varepsilon}$ and

$$
\operatorname{Pr}_{\left(g^{1}, g^{2}\right) \sim \mathcal{D}}\left[g^{2} \leq 8 n \exp \left(-\frac{1}{8} \varepsilon^{2} n\right)\right]<\frac{\varepsilon}{4 n} .
$$

Note that this is a restatement of Equation (9) applied to $\frac{\varepsilon}{4}$. It follows from the union bound that

$$
\begin{equation*}
\operatorname{Pr}_{\left(g^{1}, g^{2}\right) \sim \mathcal{D}^{n}}\left[\min _{i \in[n]}\left\{g_{i}^{2}\right\} \leq 8 n \exp \left(-\frac{1}{8} \varepsilon^{2} n\right)\right]<\frac{\varepsilon}{4} \tag{15}
\end{equation*}
$$

For any division $q \in[-1,1]^{n}$, let $X=X(q)$ be the random variable as defined in Equation (5), which is the chooser's value for pile 1 minus the chooser's value for pile 2 . Write $\mu_{X} \in[-n, n]$ for the expectation of $X$ conditioned on the divider's values. Assume we draw a list of divider values $g_{1}^{2}, g_{2}^{2}, g_{3}^{2}, \ldots, g_{n}^{2}$ such that the event in Equation (15) does not occur, i.e.,

$$
\begin{equation*}
\min _{i \in[n]}\left\{g_{i}^{2}\right\}>8 n \exp \left(-\frac{1}{8} \varepsilon^{2} n\right) \tag{16}
\end{equation*}
$$

Then, under any division,

$$
\begin{aligned}
\operatorname{Pr}\left[\left|X-\mu_{X}\right| \geq \frac{\varepsilon}{2} n\right] & \leq \operatorname{Pr}\left[\left|X-\mu_{X}\right| \geq \frac{\varepsilon}{2} n-1\right] \\
& \leq \operatorname{Pr}\left[\left|X-\mu_{X}\right| \geq \frac{\varepsilon}{4} n\right] \quad(\text { since } n \geq 4 / \varepsilon) \\
& \leq 2 \exp \left(-\frac{2\left(\frac{\varepsilon}{4} n\right)^{2}}{\sum_{i=0}^{n}(1-0)^{2}}\right) \quad \text { (Hoeffding's inequality) } \\
& =2 \exp \left(-\frac{1}{8} \varepsilon^{2} n\right) \\
& <\frac{\min _{i \in[n]}\left\{g_{i}^{2}\right\}}{4 n} \quad \text { (from Equation (16)) } \\
& \leq \frac{1}{4 n} \\
& \leq \frac{1}{n} \\
& \leq \frac{\varepsilon}{4} \quad(\text { since } n \geq 4 / \varepsilon) .
\end{aligned}
$$

There are three inequalities derived from this chain that we will use shortly:

$$
\begin{align*}
& \operatorname{Pr}\left[\left|X-\mu_{X}\right| \geq \frac{\varepsilon}{2} n\right]<\frac{1}{4 n}<\frac{1}{2}  \tag{17}\\
& \operatorname{Pr}\left[\left|X-\mu_{X}\right| \geq \frac{\varepsilon}{2} n-1\right]<\frac{\min _{i \in[n]}\left\{g_{i}^{2}\right\}}{4 n}  \tag{18}\\
& \operatorname{Pr}\left[\left|X-\mu_{X}\right| \geq \frac{\varepsilon}{2} n\right]<\frac{\varepsilon}{4} . \tag{19}
\end{align*}
$$

We claim that, still assuming Equation (16) holds, $\left|\mu_{X}\right| \leq \frac{\varepsilon}{2} n$ under any optimal division. First suppose that $\mu_{X}>\frac{\varepsilon}{2} n$. By Equation (17), with probability at least $\frac{1}{2}$, we will have $\left|X-\mu_{X}\right|<\frac{\varepsilon}{2} n$, in which case the triangle inequality implies $|X|>0$. This means the chooser will pick pile 1 with probability strictly greater than $\frac{1}{2}$, which cannot happen in an optimal division, so we have a contradiction. Next suppose that $\mu_{X}<-\frac{\varepsilon}{2} n$. Note that any optimal division $q$ has some $q_{j} \leq 0$; otherwise the chooser picks pile 1 with certainty. Take such a $j$, and consider the alternative division $q^{\prime}$ where $q_{j}^{\prime}:=q_{j}+1$ and $q_{i}^{\prime}:=q_{i}$ for all $i \neq j$. Let $X^{\prime}$ be the new difference in pile utilities under $q^{\prime}$, with mean $\mu_{X^{\prime}}$. Let $P$ and $P^{\prime}$ be the respective probabilities that the chooser picks pile 1 in $q$ and $q^{\prime}$, and let $u$ and $u^{\prime}$ be the respective expected divider utilities. Since the value of the $j^{\text {th }}$ good has expectation at most 1 , we have $\mu_{X^{\prime}} \leq \mu_{X^{\prime}}+1<-\left(\frac{\varepsilon}{2} n-1\right)$. Thus, $X^{\prime}>0$ implies $\left|X^{\prime}-\mu_{X^{\prime}}\right| \geq \frac{\varepsilon}{2} n-1$, so

$$
\begin{equation*}
P^{\prime}=\operatorname{Pr}\left[X^{\prime}>0\right] \leq \operatorname{Pr}\left[\left|X^{\prime}-\mu_{X^{\prime}}\right| \geq \frac{\varepsilon}{2} n-1\right]<\frac{\min _{i \in[n]}\left\{g_{i}^{2}\right\}}{4 n} \leq \frac{\left(q_{j}^{\prime}-q_{j}\right) g_{j}^{2}}{4 n} . \tag{21}
\end{equation*}
$$

In the middle inequality, we have applied Equation (18) to $X^{\prime}$, which is valid since the bound holds under any division. Note in particular this means that $P^{\prime} \leq \frac{1}{4}$. We may thus lower bound the utility gain of the divider from switching to $q^{\prime}$ as follows:

$$
\begin{aligned}
u^{\prime}-u & =\left(\sum_{i=1}^{n} \frac{g_{i}^{2}}{2}+\left(\frac{1}{2}-P^{\prime}\right) \sum_{i=1}^{n} q_{i}^{\prime} g_{i}^{2}\right)-\left(\sum_{i=1}^{n} \frac{g_{i}^{2}}{2}+\left(\frac{1}{2}-P\right) \sum_{i=1}^{n} q_{i} g_{i}^{2}\right) \\
& =\left(\frac{1}{2}-P^{\prime}\right) \sum_{i=1}^{n} q_{i}^{\prime} g_{i}^{2}-\left(\frac{1}{2}-P\right) \sum_{i=1}^{n} q_{i} g_{i}^{2} \\
& =\left(\frac{1}{2}-P^{\prime}\right)\left(\sum_{i=1}^{n} q_{i}^{\prime} g_{i}^{2}-\sum_{i=1}^{n} q_{i} g_{i}^{2}\right)+\left(\left(\frac{1}{2}-P^{\prime}\right)-\left(\frac{1}{2}-P\right)\right) \sum_{i=1}^{n} q_{i} g_{i}^{2} \\
& =\left(\frac{1}{2}-P^{\prime}\right)\left(q_{j}^{\prime}-q_{j}\right) g_{j}^{2}+\left(P-P^{\prime}\right) \sum_{i=1}^{n} q_{i} g_{i}^{2} \\
& \geq\left(\frac{1}{2}-P^{\prime}\right)\left(q_{j}^{\prime}-q_{j}\right) g_{j}^{2}-P^{\prime} \sum_{i=1}^{n} q_{i} g_{i}^{2} \\
& >\left(\frac{1}{4}\right)\left(q_{j}^{\prime}-q_{j}\right) g_{j}^{2}-P^{\prime} n \quad\left(\text { since } P^{\prime}<\frac{1}{4} \text { and each } q_{i} g_{i}^{2} \leq 1\right) \\
& \left.>\left(\frac{1}{4}\right) P^{\prime} \cdot 4 n-P^{\prime} n \quad \quad \text { (from Equation }(21)\right) \\
& =0
\end{aligned}
$$

Thus, we have shown that $q^{\prime}$ is a strictly better division, contradicting the assumption that $q$ was optimal. It follows that $\left|\mu_{X}\right| \leq \frac{\varepsilon}{2}$.

Therefore, we know from Equations (15) and (19) that the conditions $\left|X-\mu_{X}\right|<\frac{\varepsilon}{2} n$ and $\left|\mu_{X}\right| \leq \frac{\varepsilon}{2}$ are each violated with probability less than $\frac{\varepsilon}{4}$. By the union bound, it follows that, with probability at least $1-\frac{\varepsilon}{2}$, neither condition is violated. In this case, we know that $|X|<\varepsilon n$ by the triangle inequality. Thus, with probability at least $1-\frac{\varepsilon}{2}$, we may bound $X$ by $\varepsilon n$, whereas in the other case we may still bound $X$ by $n$. Putting these bounds together, we conclude that the chooser's expected utility is

$$
\mathbb{E}\left[u^{C}\right]=n B+\frac{\mathbb{E}[|X|]}{2} \leq n B+\frac{\left(1-\frac{\varepsilon}{2}\right) \varepsilon n+\left(\frac{\varepsilon}{2}\right) n}{2}<n(B+\varepsilon)
$$

as desired.


[^0]:    ${ }^{1}$ Throughout this paper, we assume without loss of generality that "pile 1 " is the divider's preferred pile.

[^1]:    ${ }^{2}$ Throughout this section we use the word "prior" to refer to the distribution of chooser values conditioned on the divider's values.

[^2]:    ${ }^{3}$ In general, computing volumes of polyhedra is \#P-hard [16]. The best known formula for this particular problem involves enumerating the exponentially-many vertices [17].

[^3]:    ${ }^{4}$ Technically speaking, the division output by Algorithm 1 is not guaranteed to be "close" to the globally optimal division, even though the objective values must be similar. However, for the divisions in Figures 2, 4, and 8, we did verify that the optimal division, $p^{*}$, must be close to the computed division shown in the figure, $p$, in the sense that $\left\|p^{*}-p\right\|_{\infty} \leq 0.05$. (For Figure 8, we only obtained 0.1 instead of 0.05 .) This is a close enough approximation to conclude that the properties we are claiming hold in these two examples. We computed this by re-running Algorithm 1 with an additional constraint that each $p_{i}$ be bounded away from the value in the original computed solution by 0.05 . We did this separately for each good $i$, and each direction of the deviation (bounded away from above/below). For a small enough value of the error parameter $\gamma$, we can conclude that, in each of these $2 n$ constrained optimizations, the optimal objective value decreases by more than $\gamma$. Thus, the globally optimal solution $p^{*}$ cannot respect any of these additional constraints, implying that $\left\|p^{*}-p\right\|_{\infty} \leq 0.05$.

[^4]:    ${ }^{5}$ Indeed, that fraction is zero or one, except for a single good intended to keep the piles even in value to the chooser.

[^5]:    ${ }^{6}$ Even though the expected utilities are closer at 15 goods, the probability that the chooser is better off than the divider is closer to $\frac{1}{2}$ at 13 goods than at 15 goods.

