



Playing Divide-and-Choose Given Uncertain Preferences

Citation

Tucker-Foltz, Jamie, and Richard Zeckhauser. "Playing Divide-and-Choose Given Uncertain Preferences," M-RCBG Faculty Working Paper Series 2022-01, Harvard University, Cambridge, MA, July 2022.

Permanent link

<https://nrs.harvard.edu/URN-3:HUL.INSTREPOS:37372785>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)



HARVARD Kennedy School

MOSSAVAR-RAHMANI CENTER
for Business and Government

Playing Divide-and-Choose Given Uncertain Preferences

Jamie Tucker-Foltz

Harvard University

Richard J. Zeckhauser

Harvard University

2022

M-RCBG Faculty Working Paper Series | 2022-01

Mossavar-Rahmani Center for Business & Government

Weil Hall | Harvard Kennedy School | www.mrcbg.org

The views expressed in the M-RCBG Associate Working Paper Series are those of the author(s) and do not necessarily reflect those of the Mossavar-Rahmani Center for Business & Government or of Harvard University. The papers in this series have not undergone formal review and approval; they are presented to elicit feedback and to encourage debate on important public policy challenges. Copyright belongs to the author(s). Papers may be downloaded for personal use only.

Playing Divide-and-Choose Given Uncertain Preferences

Jamie Tucker-Foltz* and Richard Zeckhauser†
Harvard University

July 13, 2022

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 private values for the goods. The prior distributions on those values are independent and common knowledge.

We characterize the structure of optimal divisions in the divide-and-choose game and show how to efficiently compute equilibria. We identify several striking differences between optimal strategies in the cases of known versus unknown preferences. Most notably, the divider has a compelling “diversification” incentive in creating the chooser’s two options. This incentive, hereto unnoticed, 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’ uncertainties about each other’s values, the number of goods being divided, and whether the divider can offer multiple alternative divisions. We prove that, when values are independently and identically distributed across players and goods, 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.

Sometimes, as with the famed cake-cutting problem, the assets are continuous. Other times, there are separate indivisible assets, as perhaps valuables from an estate. Consider cake-cutting, where the cake has vanilla filling, chocolate icing, and a cherry, each element being divisible. One player will divide the cake; the other will then choose her piece. Posit that the two players have additive preferences and that fractional portions of assets take proportional fractional values. If the players’ preferences are known, the divider 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.

*jtuckerfoltz@gmail.com

†richard_zeckhauser@hks.harvard.edu

¹Throughout this paper, we assume without loss of generality that “pile 1” is the divider’s preferred pile.

Or consider, as a more consequential example, partners in a real-estate company with 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.

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. 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 lower 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. That is, the players' values are private information, though the priors for these 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 when preferences are known, though real-world situations rarely meet that standard. In the comforting land of known preferences, the divider has a simple optimization task. He simply allocates goods based on the ascending ratio of his own value to the chooser's value until the chooser just picks pile 2. 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 and two-point discrete distributions.

This complexity arises because it is often beneficial to divide more than one good.² 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 probability that the chooser will opt for the divider's preferred pile 1. In Section 4 we investigate the divider's incentive to diversify his risk by dividing several goods (or, for indivisible assets, using a lottery to conduct such divisions). Even for a risk-neutral divider, diversification is warranted for a wide range of prior distributions. Furthermore, we show that, when the relevant goods are divisible, risk-aversion increases the extent to which the divider should diversify. These results still do not determine *which* goods should be split between the two piles. We show that simple rules can lead the divider astray.

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, however, is a method that effectively accomplishes the task for independent normal priors. That method efficiently computes divisions that yield expected utility that comes arbitrarily close to the maximal expected utility (Section 3.2). Formally, we prove that this algorithm is a fully *polynomial-time approximation scheme*. We also present an algorithm that solves the case of discrete priors, which is practical when the number of possible chooser types is small.

²With known preferences, to secure equality of the chooser's total value, at most good need be divided between the two piles.

Finally, in Section 5 we analyze the expected utilities of the divider and chooser under a range of circumstances. For example, if both players’ values are drawn from the same distribution, which player is better off *a priori*? 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. The divider’s utility increases notably if he is allowed to make multiple alternative offers to divide the goods in different manners. Surprisingly, even in the case of i.i.d. priors, the divider’s ability to make multiple offers can *decrease* the chooser’s expected utility.

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 *cake-cutting* model [20], 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 the highest.
- (ii) The protocol can be implemented by asking the players simple queries.

Extensions have also been discovered for 3 or more players [6, 2]. These methods require several rounds of dividing and choosing, but ultimately satisfy the same two properties. 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 [8, 12].

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 [11, 10, 4, 5, 9] 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; we are not interested in finding mechanisms to implement socially desirable outcomes. Rather, we ask how strategic players should actually behave to maximize their personal welfares. 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 [18]. On the theoretical front, however, much less is known. We are aware of one prior work, by Delgosha and Gohari [14], 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.

Our problem bears many similarities to the problem of optimal pricing of multiple goods [13]. In that case, the seller (the divider) must post optimal prices (a division of goods) taking into account the uncertainty over the buyer’s (chooser’s) value. We may think of the divider as offering the chooser a bundle of goods in one of the piles, hoping that the chooser takes that pile and leaves the remaining (more valuable goods to him) to the divider, just as a seller hopes that the buyer will accept a given bundle at the designated posted price. Of course, the two settings are not entirely isomorphic, as there is no money being transferred; instead the divider’s utility is affected in a more subtle way. A further difference is that the divider is constrained to offer only a single bundle of goods to the chooser, whereas a seller of multiple goods can (and often should) post several different options. Restricting the number of offers has also been considered in the literature on the optimal pricing of multiple goods [16, 17]. In Section 5.2 we consider a realistic extension of the divide-and-choose game that essentially amounts to relaxing the constraint of having only a single bundle.

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 analysis (see, for example, Weinstein and Zeckhauser [21]). In the case where the divider has complete 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 [22], as well as Brânzei, Caragiannis, Kurokawa, and Procaccia [7]. 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.

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 $[n] := \{1, 2, 3, \dots, n\}$ 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 4.2 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 independently from respective distributions \mathcal{G}_i^D and \mathcal{G}_i^C . More compactly, we denote the value vectors $g^D := (g_1^D, g_2^D, \dots, g_n^D)$ and $g^C := (g_1^C, g_2^C, \dots, g_n^C)$, which are drawn from respective distributions $\mathcal{G}^D := \mathcal{G}_1^D \times \mathcal{G}_2^D \times \dots \times \mathcal{G}_n^D$ and $\mathcal{G}^C := \mathcal{G}_1^C \times \mathcal{G}_2^C \times \dots \times \mathcal{G}_n^C$. The divide-and-choose game proceeds in three steps.

- (i) For each good i , the divider privately observes $g_i^D \sim \mathcal{G}_i^D$ and the chooser privately observes $g_i^C \sim \mathcal{G}_i^C$.
- (ii) The divider chooses a *division* of the goods, which is a vector $p = (p_1, p_2, \dots, p_n) \in [0, 1]^n$. We refer to *pile 1* as the allocation consisting of p_i of each good i and *pile 2* as the allocation consisting of $(1 - p_i)$ 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

$$u^D := \sum_{i=1}^n (1 - p_i) g_i^D$$

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

$$u^C := \sum_{i=1}^n (1 - p_i) g_i^C.$$

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 distributions \mathcal{G}_i^D and \mathcal{G}_i^C . We focus on the following distributions:

- *Normal priors*: The value for good each i is drawn from $\mathcal{N}(\mu_i, \sigma_i)$ 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 values for all goods are drawn i.i.d. from the uniform distribution on $[0, 1]$.

Throughout this paper we often assume that each g_i^D has been fixed, in which case the distribution \mathcal{G}_i^D is irrelevant. Given fixed values of g_i^D , we define the *critical ratio* of good i to be g_i^D divided by the expectation of $g_i^C \sim \mathcal{G}_i^C$.

Note that our model allows for goods with potentially negative values (i.e., “bads”). With normal priors this will happen with some nonzero probability, though in most of our examples the mean is sufficiently large to render this probability negligibly small. We only discuss critical ratios in the context where all means are positive.

We later consider extensions of this game whereby the divider may simultaneously propose multiple divisions, allowing the chooser to select whichever she wishes. We introduce notation for this setting when needed.

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 their 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 $g_1^D, g_2^D, \dots, g_n^D$. How should he use that information together with the known priors on the chooser's values to determine his optimal division? We first establish some basic, general properties of the equilibria of the divide-and-choose game. We then present algorithms to determine the divider's optimal division for both normal priors and discrete priors.

3.1 The general structure of optimal divisions

We begin with the following observation, which, 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 ,*

$$\mathbb{E}_{g^C \sim \mathcal{G}^C} [u^D] \geq \frac{1}{2} \sum_{i=1}^n g_i^D.$$

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

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

so at least one of the options must yield at least this utility. To prove (ii), observe that the division $p = (\frac{1}{2}, \frac{1}{2}, \dots, \frac{1}{2})$ 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. \square

We refer to the quantities

$$\frac{1}{2} \sum_{i=1}^n g_i^D \quad \text{and} \quad \frac{1}{2} \sum_{i=1}^n g_i^C$$

as the respective *baseline* 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

$$q_i := 2p_i - 1 = p_i - (1 - p_i), \quad (1)$$

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

$$p_i = \frac{q_i}{2} + \frac{1}{2}. \quad (2)$$

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:*

$$\sum_{i=1}^n q_i g_i^D \geq 0. \quad (3)$$

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

$$P := \Pr_{g^C \sim \mathcal{G}^C} \left[\sum_{i=1}^n q_i g_i^C \geq 0 \right] \leq \frac{1}{2}. \quad (4)$$

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

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

$$\begin{aligned} \mathbb{E}[u^D] &= P \left(\sum_{i=1}^n (1 - p_i) g_i^D \right) + (1 - P) \left(\sum_{i=1}^n p_i g_i^D \right) \\ &= P \left(\sum_{i=1}^n \left(1 - \frac{q_i}{2} - \frac{1}{2} \right) g_i^D \right) + (1 - P) \left(\sum_{i=1}^n \left(\frac{q_i}{2} + \frac{1}{2} \right) g_i^D \right) \\ &= P \sum_{i=1}^n \frac{g_i^D}{2} - P \sum_{i=1}^n \frac{g_i^D q_i}{2} + \sum_{i=1}^n \frac{g_i^D q_i}{2} + \sum_{i=1}^n \frac{g_i^D}{2} - P \sum_{i=1}^n \frac{g_i^D q_i}{2} - P \sum_{i=1}^n \frac{g_i^D}{2} \\ &= \sum_{i=1}^n \frac{g_i^D}{2} + \left(\frac{1}{2} - P \right) \sum_{i=1}^n q_i g_i^D. \end{aligned} \quad (5)$$

The optimal utility is attained by maximizing Equation (5) over the variables $q_1, q_2, \dots, q_i \in [-1, 1]$. Let u^* denote the supremum of this optimal utility, and consider a sequence of divisions q^1, q^2, q^3, \dots 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 (5) 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 (5) is nonnegative by (3), we have

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

so the divider is no better off than his baseline utility. Therefore, the divider is at least equally well-off setting $p = (\frac{1}{2}, \frac{1}{2}, \dots, \frac{1}{2})$, 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 (5) is nonzero if and only if the inequalities in both Equations (3) and (4) are strict. \square

3.2 Normal priors

In the case where the divider has normal priors over the chooser's value, we can characterize exactly when the inequalities from Lemma 3.2 will be strict.

Proposition 3.3. *Suppose each \mathcal{G}_i^C is a normal distribution with positive mean, and suppose all divider values are positive. Then the optimal division yields baseline utility if and only if all goods have the same critical ratios.*

Proof. Suppose first that all goods have the same critical ratio r . Then observe that, for any division q ,

$$\begin{aligned} P > \frac{1}{2} &\iff \Pr \left[\sum_{i=1}^n q_i g_i^C > 0 \right] > \frac{1}{2} \\ &\iff \mathbb{E} \left[\sum_{i=1}^n q_i g_i^C \right] > 0 \quad (\text{since } \sum_{i=1}^n q_i g_i^C \text{ is normally distributed}) \\ &\iff \mathbb{E} \left[\sum_{i=1}^n q_i g_i^D \frac{g_i^C}{g_i^D} \right] > 0 \\ &\iff \sum_{i=1}^n q_i g_i^D \frac{\mathbb{E}[g_i^C]}{g_i^D} > 0 \\ &\iff r \sum_{i=1}^n q_i g_i^D > 0 \\ &\iff \sum_{i=1}^n q_i g_i^D > 0 \quad (\text{since } r \text{ is positive by assumption}). \end{aligned}$$

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-baseline utility by the final statement of Lemma 3.2.

Now suppose that two goods j and k have different critical ratios, i.e., \mathcal{G}_j^C has mean μ_j and \mathcal{G}_k^C has mean μ_k such that,

$$\frac{g_j^D}{\mu_j} > \frac{g_k^D}{\mu_k}.$$

Let $\alpha := \max\{g_j^D + \mu_j, g_k^D + \mu_k\}$, let q be the division such that

$$\begin{aligned} q_j &:= \frac{g_k^D + \mu_k}{\alpha} \\ q_k &:= -\frac{g_j^D + \mu_j}{\alpha} \\ q_i &:= 0 \quad (\text{for all } i \in [n] \setminus \{j, k\}). \end{aligned}$$

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 α ensures $|q_j|, |q_k| \leq 1$). 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 $\sum_{i=1}^n q_i g_i^C$, which determines whether the chooser picks pile 1, is normally distributed. It 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 that

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

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

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 baseline; see Proposition 4.1.

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 may rewrite the divider’s optimization problem as maximizing

$$\mathbb{E}[u^D] = \sum_{i=1}^n \frac{g_i^D}{2} (P(1 - q_i) + (1 - P)(1 + q_i))$$

over the variables $q_1, q_2, \dots, q_n \in [-1, 1]$. This is *almost* a linear program. Indeed, 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, Φ denotes the standard normal cumulative distribution function.

Algorithm 1: Computes an approximately optimal division given the divider's values and independent normal priors for the chooser's values.

Input: Divider values $g_1^D, g_2^D, \dots, g_n^D$ (not all zero), prior means $\mu_1, \mu_2, \dots, \mu_n$, corresponding standard deviations $\sigma_1, \sigma_2, \dots, \sigma_n$, and an additive error bound γ for the divider's optimal utility.

Output: An approximately optimal division p_1, p_2, \dots, p_n .

```

1  $\delta \leftarrow \frac{\gamma}{\sum_{i=1}^n |g_i^D|}$ 
2  $P \leftarrow \frac{1}{2}$ 
3  $u \leftarrow -\infty$ 
4 while  $P > 0$  do
5    $u_P, q_1, q_2, \dots, q_i \leftarrow$  optimal solution to the following program  $\mathcal{C}_P$ :
      maximize    $u_P = \sum_{i=1}^n \frac{g_i^D}{2} (P(1 - q_i) + (1 - P)(1 + q_i))$ 
      subject to  $-1 \leq q_i \leq 1$  for all  $1 \leq i \leq n$ ,
                   $\sum_{i=1}^n g_i^D q_i \geq 0$ ,
                   $\sum_{i=1}^n \mu_i q_i \leq \Phi^{-1}(P) \sqrt{\sum_{i=1}^n \sigma_i^2 q_i^2}$ 
6   if  $u_P > u$  then
7      $u \leftarrow u_P$ 
8     for  $i \leftarrow 1, 2, \dots, n$  do
9        $p_i \leftarrow \frac{q_i + 1}{2}$ 
10    end
11   $P \leftarrow P - \delta$ 
12 end
13 return  $(p_1, p_2, \dots, p_n)$ 

```

Lemma 3.4. *Algorithm 1 runs in polynomial time in the values of n and $\frac{\sum_{i=1}^n |g_i^D|}{\gamma}$.*

Proof. Observe that there are at most

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

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 $q := (q_1, q_2, \dots, q_n) \in S$, and let $\tilde{q} := (\sigma_1 q_1, \sigma_2 q_2, \dots, \sigma_n q_n)$. Then, for any positive real number c , the scaled vector $cq = (\frac{q_1}{c}, \frac{q_2}{c}, \dots, \frac{q_n}{c})$

lies in S as well, since

$$\begin{aligned}
\sum_{i=1}^n \mu_i(cq_i) &= c \cdot \sum_{i=1}^n \mu_i q_i \\
&\leq c \cdot \Phi^{-1}(P) \sqrt{\sum_{i=1}^n \sigma_i^2 q_i^2} \quad (\text{since } q \in S) \\
&= \Phi^{-1}(P) c \|\tilde{q}\|_2 \quad (\text{where } \|\cdot\|_2 \text{ denotes the Euclidean norm}) \\
&= \Phi^{-1}(P) \|c\tilde{q}\|_2 \quad (\text{by linearity of the Euclidean norm}) \\
&= \Phi^{-1}(P) \sqrt{\sum_{i=1}^n \sigma_i^2 (cq_i)^2}.
\end{aligned}$$

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

We remark that, beyond being efficient in theory, this algorithm proves to be fast in practice. We implemented this algorithm using Gurobi [15] 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 γ of the optimal divider utility.*

The proof is long and technical, so is deferred to Appendix A.

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 polynomial-time approximation scheme with approximation parameter ε .*

Proof. Given this choice of γ , it follows from Lemma 3.4 and the fact that $|g_i^D| = g_i^D$ 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,

$$2u^* \geq \sum_{i=1}^n g_i^D.$$

Combining this inequality with Lemma 3.5, we have

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

as desired. \square




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 $\mathcal{G}_1^C = \mathcal{N}(5, 1)$, $\mathcal{G}_2^C = \mathcal{N}(9.5, 1)$, $\mathcal{G}_3^C = \mathcal{N}(13.6, 9.8)$, $\mathcal{G}_4^C = \mathcal{N}(95, 169)$.

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, this is not always the case, 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 $\mathcal{G}_1^C = \mathcal{N}(100, 1)$, $\mathcal{G}_2^C = \mathcal{N}(100, 1)$, $\mathcal{G}_3^C = \mathcal{N}(100, 65)$, 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 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. \square

3.3 Discrete priors

Now we turn to the case where the prior \mathcal{G}^C is an arbitrary distribution supported on a finite set $\{x_1, x_2, \dots, x_\ell\} \subseteq \mathbb{R}^n$. For each $1 \leq j \leq \ell$, we write r_j for the probability that $g^C = x_j$. In this case, we can exactly solve for the optimal division, but the algorithm is not efficient. The main idea 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 ℓ is small. This algorithm has the additional advantage that it works even when players' values are correlated across multiple goods.

Algorithm 2: Computes an optimal division given the divider's values and an arbitrary discrete prior for the chooser's values.

Input: Divider values $g_1^D, g_2^D, \dots, g_n^D$ (not all zero), possible chooser value vectors $x_1, x_2, \dots, x_\ell \in \mathbb{R}^n$, and corresponding probabilities r_1, r_2, \dots, r_ℓ .

Output: An optimal division p_1, p_2, \dots, p_n .

```

1  $u \leftarrow -\infty$ 
2 for  $S \subseteq [\ell]$  do
3    $P \leftarrow \sum_{j \in S} r_j$ 
4   if  $P \leq \frac{1}{2}$  then
5      $u', q_1, q_2, \dots, q_i \leftarrow$  optimal solution to the following linear program  $\mathcal{C}_S$ :
           
$$\begin{array}{ll} \text{maximize} & u' = \sum_{i=1}^n \frac{g_i^D}{2} (P(1 - q_i) + (1 - P)(1 + q_i)) \\ \text{subject to} & -1 \leq q_i \leq 1 \quad \text{for all } 1 \leq i \leq n, \\ & \sum_{i=1}^n (x_j)_i q_i \leq 0 \quad \text{for all } j \in [\ell] \setminus S \end{array}$$

6     if  $u' > u$  then
7        $u \leftarrow u'$ 
8       for  $i \leftarrow 1, 2, \dots, n$  do
9          $p_i \leftarrow \frac{q_i + 1}{2}$ 
10      end
11    end
12 end
13 return  $(p_1, p_2, \dots, p_n)$ 

```

Proposition 3.8. *Algorithm 2 finds a division yielding optimal divider utility in time $\text{poly}(n)$ for any constant ℓ .*

Proof. Let $p^* = (p_1^*, p_2^*, \dots, p_n^*)$ 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' 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. \square

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 framework 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 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. 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 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, but 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. Consider goods 1 and 2 (of many others), for which the chooser’s values are either 0 or 1, equally likely. 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 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 4.2, 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 respectively Algorithm 1 and Algorithm 2. Here we provide a qualitative explanation of how goods are optimally allocated between the two piles.




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 $g_1^D = 101$, $g_2^D = 102$, $g_3^D = 103$, $g_4^D = 104$, $g_5^D = 105$, $g_6^D = 200$, and the chooser’s prior for each good is $\mathcal{G}_i^C = \mathcal{N}(10, 1)$. The optimal value of P is 0.034.

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 1.³

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.

Although the first five goods are diversified, the ones the divider values more tilt more toward pile 1 than 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 .⁴

³Technically 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, 3, and 6, we did verify that the optimal division, p^* , must be close to the computed division shown in the figure, p , in the sense that $\|p^* - p\|_\infty \leq 0.05$. (For Figure 6, 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 γ , we can conclude that, in each of these $2n$ constrained optimizations, the optimal objective value decreases by more than γ . Thus, the globally optimal solution p^* cannot respect any of these additional constraints, implying that $\|p^* - p\|_\infty \leq 0.05$.

⁴Indeed, that fraction is zero or one, except for a single good intended to keep the piles even in value to the chooser.

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 this may happen even with normal priors with *identical* variances.




Figure 3: 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 $\mathcal{G}_1^C = \mathcal{N}(100, 5)$, $\mathcal{G}_2^C = \mathcal{N}(198, 5)$, $\mathcal{G}_3^C = \mathcal{N}(100, 5)$. The optimal value of P is 0.005.

Figure 3 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, 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.

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 dramatically breaks this result.

A further departure from the case of certain preferences is that the divider can achieve a higher-than-baseline 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 multi-point) discrete priors, despite such identical values across goods, it is feasible for the divider to beat the baseline.

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 $\mathcal{G}_i^C = \mathcal{D}$, there exists a division of goods yielding utility higher than the divider's baseline.*

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

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

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

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.

Lemma 4.2. *If the divider can achieve a higher-than-baseline utility, he will always leave one good entirely in one of the two piles.*

Proof. Let q be an optimal division from Lemma 3.2. If q achieves a higher-than-baseline utility, then there must be some good i such that $|q_i| > 0$. Let i^* be a good maximizing $|q_{i^*}|$. If $|q_{i^*}| = 1$, 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 $|q_{i^*}|$. Since

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

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. \square

We 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 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 change if he is averse to risk? Thus, suppose the divider is an expected utility maximizer with utility $u^D = f(v^D)$, 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. Fix divider values and let f be a concave utility function. The chooser is no less likely to choose pile 1 under an optimal division assuming the divider is risk-neutral than under an optimal division assuming the divider has utility function f .

Proof. Let p be the optimal division by a risk-neutral divider, and let p' 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', P' denote the divider's value for pile 1 and the probability the chooser picks pile 1 according to p and p' , respectively. Suppose toward a contradiction that $P < P'$. Then we must have $v < v'$, for otherwise p would be a strictly better division than p' 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' , it must be that

$$(1 - P')v' + P'(T - v') < (1 - P)v + P(T - v).$$

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




Figure 4: Illustration accompanying the proof of Theorem 4.3.

Thus, values are ordered exactly as shown in Figure 4, where the x -coordinate of points A and C are the respective expected divider values of p and p' . 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' , 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' . \square

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 5 plots the optimal expected divider utility given any value of P , under 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.




Figure 5: 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^* .

Theorem 4.3 does not hold when goods are indivisible. 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 $\frac{8}{3}$.

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 $\mathcal{N}(1, 0.2)$. Figure 6 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 $\mathcal{N}(1, 0.2)$. As one can see, risk-aversion leads to more goods being split between the two piles, and generally split more equally. Still, one good is always left undivided, as Lemma 4.2 holds with the same proof.




Figure 6: 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 $\mathcal{G}_i^C := \mathcal{N}(1, 0.2)$, and the divider’s values were also sampled from $\mathcal{G}_i^D := \mathcal{N}(1, 0.2)$. 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).

5 Welfares of the players

We now turn to analyze the expected welfares of the players. Knowing their expected welfares is important if our concern is 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. We begin by studying the expected welfares of the two players, both theoretically and empirically. We then investigate how these welfares change in a realistic setting that allows for further negotiation and/or richer strategies beyond the rules of the original divide-and-choose game.

5.1 Is it better to divide or choose?

The cake-cutting literature suggests the chooser is better off, as the divider is compelled to divide goods roughly evenly, while the chooser can get a more favorable outcome because she knows her own preferences. However, if the divider has strong knowledge of the chooser’s preferences, then

the divider can exploit this knowledge [22, 7]. Indeed, as Nicoló and Yu [19] 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: under mild assumptions on the distribution of players’ values, the chooser is favored when n is small; the divider is favored when n is large.

Theorem 5.1. *Let \mathcal{D} be a probability distribution such that:*

- *\mathcal{D} is supported on at least two distinct values.*
- *The support of \mathcal{D} is bounded and nonnegative.*
- *The expectation of the minimum of n i.i.d. draws from \mathcal{D} is either bounded away from zero, or vanishes subexponentially as $n \rightarrow \infty$.*

Suppose that, for each good i , $\mathcal{G}_i^D = \mathcal{G}_i^C = \mathcal{D}$. Then, in any equilibrium, the following hold.

(i) *For $n = 2$, the chooser is ex ante strictly better off:*

$$\mathbb{E}_{g^D, g^C \sim \mathcal{D}^2} [u^D] < \mathbb{E}_{g^D, g^C \sim \mathcal{D}^2} [u^C]$$

(ii) *For all sufficiently large n , the divider is ex ante strictly better off:*

$$\mathbb{E}_{g^D, g^C \sim \mathcal{D}^2} [u^D] > \mathbb{E}_{g^D, g^C \sim \mathcal{D}^2} [u^C]$$

The assumptions on \mathcal{D} are tight in several ways. If \mathcal{D} is supported on only a single value, then clearly the divider and chooser are equally well-off for any number of goods: the divider will always choose a division p such that $\sum_{i=1}^n p_i = \frac{n}{2}$ and both players will receive their equivalent baseline utilities. The theorem also fails for some distributions with negative values. For example, if \mathcal{D} is the uniform distribution on $[-1, 1]$, one can see that $P = \frac{1}{2}$ in any division, so the divider cannot surpass his baseline utility for any number of goods. However, for large n , the chooser will certainly exceed her baseline utility (assuming the divider breaks his indifference between divisions in a way that benefits the chooser).

Proof of Theorem 5.1 (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 which is not divided, 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 $p^D := f(\{g_1^D, g_2^D\}) > \frac{1}{2}$ fraction of good $\arg \max_i (g_i^D)$ in pile 1, with the rest of that good and all of the other good in pile 2. Analogously, let $p^C := f(\{g_1^C, g_2^C\})$, i.e., the amount of the chooser's preferred good that would have gone into pile 1 if the chooser had 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 if the roles had 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. If she had instead been the divider, then 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, if she had 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. Thus, he must have strictly preferred a p^C -fraction of his favorite good. Moreover, if the chooser had been the divider she would have received her least-preferred pile, obtaining at most baseline utility. Since the chooser always can get at least her baseline, the chooser is weakly better off.

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 baseline utility and the chooser exceeds baseline utility, 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. \square

The proof of statement (ii) is technical; see Appendix B. We present an intuitive argument here. 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 highest-valued 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 baseline utility from each pile, i.e., the sum of a random subset of half of her values. For any distribution satisfying the hypotheses of Theorem 5.1, the expectation of this sum is strictly less than the expectation of the sum of the top-half of values.

As an illustration of Theorem 5.1, take \mathcal{D} to be the uniform distribution on $[0, 1]$. 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.

Equilibrium utility per $U[0, 1]$ good	$n = 2$	$n \rightarrow \infty$
Baseline	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).

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 7 plots estimated utilities per good for an empirical experiment where all values are drawn from $\mathcal{N}(1, 0.2)$. The crossover appears to be around 15 goods.




Figure 7: Estimated utilities from repeated trials of the following experiment. We first draw divider values $g_1^D, g_2^D, \dots, g_n^D$ i.i.d. from $\mathcal{N}(1, 0.2)$, then compute the optimal division p with respect to those values, with $\mathcal{G}_i^C = \mathcal{N}(1, 0.2)$ for each good i as well. We compute the divider's and chooser's expected utility exactly with respect to p and $g_1^D, g_2^D, \dots, g_n^D$, where the expectation is taken over the unknown values $g_1^C, g_2^C, \dots, g_n^C$. All utilities are averaged, then normalized by dividing by the number of goods n .

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 8 plots random samples of the values of 13 goods drawn from the same distribution, $\mathcal{N}(1, 0.2)$, as in Figure 7. We chose the number of goods to be 13 since that is near the crossover point in Figure 7, where the two roles have similar ex ante utilities.⁵ 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.




Figure 8: 200 random samples of values for 13 goods drawn i.i.d. from $\mathcal{N}(1, 0.2)$, 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 $\mathcal{N}(1, 0.2)$.

5.2 The effect of multiple offers

In real life, players of an economic game sometimes negotiate to change the rules. We now consider a modification to the game where the divider can simultaneously offer multiple proposed divisions to the chooser, not merely the original two options.

⁵Even 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.

This modification can only help the divider. The divider can always use his standard divide-and-choose offer of pile 1 versus pile 2, but can add other offers, say piles 3 and 4, where the divider prefers the complements of piles 3 and 4 to pile 1, and where the chooser might prefer pile 3 and/or pile 4 more than pile 2. A chooser type who would choose pile 1 over pile 2 if those were the only two alternatives can be lured to pile 3 or 4 instead. That would happen, for example, if the chooser's values were: pile 4 \prec pile 2 \prec pile 1 \prec pile 3. In the original game, the divider would get stuck with pile 2, but with multiple offers he would instead get the complement of pile 3.

Formally, we model the *divide-and-choose game with multiple offers* as follows. In addition to a division $p \in [0, 1]^n$, the divider also selects an arbitrarily large set of alternatives $A \subseteq [0, 1]^n$. The chooser is then allowed to pick pile 1 or pile 2, in which case allocations and payoffs are defined by p in exactly the same way as in the original divide-and-choose game, or she may instead choose some alternative $a \in A$, in which case she receives a $1 - a_i$ fraction of each good i , with the divider receiving an a_i fraction of each good i .

This extended game is a plausible model of realistic divide-and-choose games with multiple offers. Is it still *fair*? Proportionality (Lemma 3.1) still holds, as neither party is required to utilize any of the alternative divisions, so both parties can still guarantee receiving half of their respective total utilities. However, the following results suggest that there may be a tinge of unfairness.

Proposition 5.2. *In the divide-and-choose game with multiple offers, there is always an optimal divider strategy using the division $p = (\frac{1}{2}, \frac{1}{2}, \dots, \frac{1}{2})$.*

Proof. Let (p, A) be a division and assume without loss of generality that the divider weakly prefers pile 1. Consider the alternative division (p', A') where $p' := (\frac{1}{2}, \frac{1}{2}, \dots, \frac{1}{2})$ and $A' := A \cup p$. For any realization of the chooser's preferences where she would have originally picked an option other than pile 1, she will still pick that option, since the only new option that has been introduced only give the chooser her baseline utility. On the other hand, for any realization of the chooser's preferences where she would have originally picked pile 1 (in which case the divider received pile 2), she must now pick one of the two new even piles, again giving both parties baseline utility, or some other option $a \in A$. By the assumption that the divider originally preferred pile 1, the divider's baseline utility is greater than his utility from the old pile 2. Also, the divider weakly prefers any $a \in A$ to the old pile 2, for otherwise he should have never offered a in the first place. Therefore, the divider is weakly better off with (p', A') in all cases. \square

Conceptually, this result says that the divider should optimally rely exclusively on his ability to make counteroffers, not even worrying whether an initial division entices the chooser to pick pile 2. Under these favorable conditions, the divider completely eliminates his risk that he will fail to achieve his baseline utility. That assurance enables him to extract more of the value. As the following theorem shows, this value may come at the expense of the chooser.

Theorem 5.3. *In comparison to the ordinary divide-and-choose game, in the divide-and-choose game with multiple offers, the divider is always weakly better off and the chooser may be strictly worse off. This can happen even in ex ante expectation, in a game with two goods where, for each good i , $\mathcal{G}_i^D = \mathcal{G}_i^C = \mathcal{D}$ for a common distribution \mathcal{D} .*

Proof. The divider is weakly better off because any division in the original divide-and-choose game is also valid in the divide-and-choose game with multiple offers, yielding the same utility.

Now suppose each good is valued at 1 or 2 by each player, independently, each with probability $\frac{1}{2}$. To analyze ex ante expected payoffs in each version of the game, there are two equally likely cases to consider.

When the divider draws the same value for each good, the set of optimal divisions without multiple offers consists of all divisions that put an equal amount of the two goods in each pile, and no matter what, the divider gets utility equal to his common value for each good. At the extremes, he could divide both goods evenly, in which case the chooser receives expected utility 1.5, or he could put one good in one pile and the other good in the other pile, in which case the chooser gets utility 1.75. With multiple offers, the optimal strategy is to offer an even division of both goods, or 0.75 of either good, in which case the divider gets utility 1.6875 but the chooser still gets utility 1.5. Thus, with the introduction of multiple offers, the chooser’s utility has either decreased or stayed the same (and with our technical assumption that indifference was previously broken in favor of the chooser, the chooser’s utility has decreased).

Now suppose the divider draws different values for each good. Without loss of generality, assume $g_1^D = 2$ and $g_2^D = 1$. Then the optimal division puts all of good 1 in pile 1 and all of good 2 in pile 2, so the chooser picks pile 2 with probability 0.75. Thus, the divider and chooser both receive utility 1.75. With multiple offers, the divider can essentially remove the option to take pile 1 in the case where the chooser strictly prefers it, by choosing $p = (\frac{1}{2}, \frac{1}{2})$ and $A = \{(1, 0)\}$. Thus, in this case, the divider’s utility strictly increases and the chooser’s utility strictly decreases, and in all other cases utilities remain unchanged. In total, with the introduction of multiple offers, if the divider has different values for each good, his expected utility has increased to 1.875 and the chooser’s expected utility has decreased to 1.625.

Thus, the chooser is, overall, strictly worse off when the divider makes multiple offers. □

Stepping back from our model and into the real world, the conceptual takeaway from these results is that the divider has a considerable advantage in making proposals, and should try to press this advantage whenever possible. The chooser, on the other hand, could sometimes benefit from committing to playing only the original game. For instance, she could tell the divider upfront that she will not consider any additional offers beyond the basic division into pile 1 and pile 2. This threat may not be credible in practice.

6 Conclusions

How can the two of us appropriately divide a collection of assets when neither of us will honestly reveal our private value to the other? That is a question that is asked by siblings, divorcing adults, and business partners, among others. One answer, which has been put into use by myriad duos, is to use the divide-and-choose method. With it, one player sets two piles, and the other chooses her preferred pile.

Many siblings discover this method on their own; game theorists followed with their cake cutting algorithms. Others are counselled into it, by books and by advisors. When the players’ preferences are known, effective procedures have been well documented. However, uncertainty about each other’s preferences is a prime feature in many divide-and-choose contexts. Yet the subject of how to divide in such contexts is effectively unstudied.

Our investigation of optimal division follows a Bayesian path; it is assumed that priors on value are common knowledge. In most analyses, the players are assumed to be symmetrically situated, hence priors on their values are identical.

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, is in fact critical. All the goods in one pile have a higher ratio than 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, assignment by critical ratios may be violated. That is, probability density in those lotteries may be non-monotonic in the ratios. Moreover, up to $n - 1$ out of n goods might optimally be divided by lottery.

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 if there is substantial variability in actual values. That is why diversification comes in. Dividing multiple goods between the two piles significantly reduces the variability in value the chooser receives in either pile. Hence, dividing in this manner reduces the likelihood that the divider receives his bad outcome, namely that the chooser picks pile 1.

Examples are developed with normal priors and two-point discrete priors. Intriguing, non-intuitive results emerge. With normally-distributed priors, if the critical ratios for the goods are all identical, the divider can achieve no more than his baseline utility, what he would get if each good were assigned by a coin flip. Yet, for more general distributions, even when goods are identically distributed and equally valued, the divider can construct bizarre, symmetry-breaking divisions that beat his baseline utility.

Computing the divider’s optimal division is easy in the known-preferences case. It is surprisingly difficult with preferences unknown. We already mentioned the annoying lack of monotonicity in critical ratios. Worse still, multiple local optima are commonly encountered. Despite these challenges, we are pleased to identify Algorithm 1, which computes divisions for normal priors in polynomial time that approximate the optimal utility to within arbitrary precision. Algorithm 2 handles arbitrary (discrete) priors when the number of possible chooser types is small, getting exact results. We suspect that exactly computing optimal divisions in either of these settings is NP-Hard. Characterizing the precise computational complexity of these problems is an intriguing question for future work.

Computational challenges plague even what appear to be straightforward divide-and-choose situations. Fortunately, some intriguing results prove intuitive once analyzed. A salient example is our result that the chooser is advantaged when the number of goods is small, but the advantage tips to the divider after a certain threshold number of goods.

In the real world, players have the potential to break the rules. Rather than making two offers to the chooser, the divider might offer her an array of alternative options. We show that the divider benefits by making multiple offers. Moreover, the chooser might be better off in expectation by insisting on just two complementary offers.

Some of the results above are predominantly of academic interest, such as those relating to the challenges in finding optimal divisions. Many others have direct implications for real-world divide-and-choose situations. They include the importance of diversification, the tip from chooser advantage to divider advantage as the number of goods increases, and the divider’s advantage and the chooser’s disadvantage when the divider makes multiple alternative offers.

A promising direction for future work would be to investigate approximately optimal divisions in terms of their structure. For instance, can it be shown that the intuitive rule of concentrating all goods with the highest critical ratios in pile 1 yields a nontrivial constant-factor approximation to the optimal expected divider utility? In the case where all goods have positive values, a “trivial” constant factor is 2, since utility is always bounded between the baseline utility and twice the baseline utility, which is the sum of the values of all goods. Does the divider actually obtain a smaller factor than 2 with this strategy? Does he obtain *any* constant-factor approximation to his

optimal gains beyond his baseline utility?

Most of our analyses assume that the players are risk-neutral. However, risk aversion exists and can be important when dividing a major asset, such as an estate or a business. We analyze how risk aversion should affect the divider's allocation, both in the context of deterministic divisions of divisible goods and randomized divisions of indivisible goods. In the latter context, how the divider should adjust when the chooser is risk-averse remains an open question. How can the divider exploit the chooser's risk aversion?

Besides the divider making multiple offers, there are other realistic extensions of the game that could be analyzed in a similar Bayesian framework. For instance, what if the basic division can be renegotiated *after* the game has ended? This introduces complicated information flows, possibly even requiring mixed strategies at equilibrium. For instance, the chooser will probably be able to confidently infer from the division which pile the divider prefers, because, for instance, it might be less valuable in expectation given the prior for the chooser's value. It might then be optimal for her to bluff and take that pile, even if she prefers the other one. If the bluff is not detected, her negotiating position will have improved. Of course, a divider that expects this behavior could then flip the piles, tricking the chooser into choosing a small pile with goods that the divider does not want. And there would be further counters, and still further counters.

A further extension would be to investigate scenarios in which values of goods are correlated between the players, as might be the case for commodities like valuable works of art. Playing the divide-and-choose game in this widely-applicable context does not appear to have been studied at all.

Another setting we have not considered in this paper is one where the roles are endogenously chosen. Not infrequently, players begin the divide-and-choose process by allowing for self-assignment. For example, in a partnership buy-sell arrangement, usually either partner can propose a division, with the other player then forced to choose. Frequently some condition (such as elapsed time) must be met before a player – the self-appointed divider – can make a proposal. Usually the proposal is that one player keeps most of the assets and pays a price to the other. This fits into our model as a divide-and-choose problem with two goods, the asset and money. Given our result that the preferred role might depend on the private values, one might expect a complex interplay between choosing roles, inferring values, and leveraging these inferences to construct optimal divisions.

Finally, there are various other related games we expect will yield to many of the techniques employed here. Another common way to divide up assets is the alternating-choice (or “round-robin”) method. It is used to pick sports teams, in contexts from playgrounds to professional sports leagues, and by children dividing their parents' household goods. A key ingredient of the alternating-choice method is a prime divide-and-choose concern: how great is the value to me relative to the distribution of value to the other player?

Take-it-or-leave-it (TIOLI) offers are also widely encountered, as when a firm makes one of its standardized job offers, or when a multi-service Internet and TV offer is sold at a fixed price. Legislation that must be approved or vetoed by a chief executive is also TIOLI. Pile 2 in these cases can be thought of as the status quo. The main difference between TIOLI and divide-and-choose is that, with TIOLI, both players receive the same outcome, and the total sum of piles is no longer exogenous.

Allocation systems, across a broad array from divide-and-choose to the market, involve players acting strategically to maximize their take given the rules and what they know of the other player's value. In the allocation of academic attention, we hope to have shown, divide-and-choose strategy has received a smaller pile than it deserves.

Acknowledgments

We are deeply grateful to Yannai Gonczarowski for his careful reading of an earlier draft, with several suggestions that improved both the quality and the clarity of our results. We are also very thankful to Justin Chan for numerous suggestions to improve readability.

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 are those of the authors and do not necessarily reflect the views of the National Science Foundation or the Mossavar-Rahmani Center.

References

- [1] Georgios Amanatidis, Georgios Birmipas, 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] Steven J. Brams and Alan D. Taylor. An envy-free cake division protocol. *The American Mathematical Monthly*, 102(1):9–18, 1995.
- [7] 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.
- [8] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [9] 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.
- [10] 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.

- [11] 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.
- [12] 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.
- [13] Constantinos Daskalakis. Multi-item auctions defying intuition? *SIGecom Exch.*, 14(1):41–75, 2015.
- [14] Payam Delgosha and Amin Gohari. Information theoretic cutting of a cake. *IEEE Transactions on Information Theory*, 63(11):6950–6978, 2017.
- [15] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2021.
- [16] Sergiu Hart and Noam Nisan. Approximate revenue maximization with multiple items. *Journal of Economic Theory*, 172:313–347, 2017.
- [17] Sergiu Hart and Noam Nisan. Selling multiple correlated goods: Revenue maximization and menu-size complexity. *Journal of Economic Theory*, 183:991–1029, 2019.
- [18] 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.
- [19] Antonio Nicolò and Yan Yu. Strategic divide and choose. *Games and Economic Behavior*, 64(1):268–289, 2008.
- [20] Ariel D. Procaccia. Cake cutting: not just child’s play. *Communications of the ACM*, 56(7):78–87, 2013.
- [21] Milton Weinstein and Richard Zeckhauser. Critical ratios and efficient allocation. *Journal of Public Economics*, 2(2):147–157, 1973.
- [22] H. Peyton Young. *Equity - in theory and practice*. Princeton University Press, 1995.

Appendix

A Proof of Lemma 3.5

Observe that each division p_1, p_2, \dots, 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 first constraint of \mathcal{C}_P . We claim that, on each iteration of the main loop, the division p_1, p_2, \dots, p_n computed by Algorithm 1 achieves an interim expected utility of $\mathbb{E}[u^D] \geq u_P$. (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

$$\sum_{i=1}^n g_i^C p_i \geq \sum_{i=1}^n g_i^C (1 - p_i).$$

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 μ_i and variance σ_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, \dots, q_n on each iteration of the main loop to satisfy the third constraint, we know that this probability is at most P . Therefore,

$$\begin{aligned} \mathbb{E}[u^D] &= \Pr[\text{chooser picks pile 1}] \sum_{i=1}^n g_i^D (1 - p_i) + \Pr[\text{chooser picks pile 2}] \sum_{i=1}^n g_i^D p_i \\ &= \sum_{i=1}^n g_i^D p_i + \Pr[\text{chooser picks pile 1}] \left(\sum_{i=1}^n g_i^D (1 - p_i) - \sum_{i=1}^n g_i^D p_i \right) \\ &\geq \sum_{i=1}^n g_i^D p_i + P \left(\sum_{i=1}^n g_i^D (1 - p_i) - \sum_{i=1}^n g_i^D p_i \right) \\ &\quad \text{(since the second constraint ensures the term in parentheses is nonpositive)} \\ &= \sum_{i=1}^n g_i^D (P(1 - p_i) + (1 - P)p_i) \\ &= \sum_{i=1}^n g_i^D \left(P \left(1 - \left(\frac{q_i}{2} + \frac{1}{2} \right) \right) + (1 - P) \left(\frac{q_i}{2} + \frac{1}{2} \right) \right) \\ &= \sum_{i=1}^n \frac{g_i^D}{2} (P(1 - q_i) + (1 - P)(1 + q_i)) \\ &= u_P. \end{aligned}$$

Thus, the claim is proved.

Let $(p_1^*, p_2^*, \dots, p_n^*)$ denote an optimal division from Lemma 3.2, yielding interim expected utility u^* , and let $(q_1^*, q_2^*, \dots, q_n^*)$ 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 (q_1, q_2, \dots, q_n) denote the optimal solution to \mathcal{C}_P on this iteration, with optimal value u_P , and let (p_1, p_2, \dots, p_n) denote the corresponding division.

Observe that $(q_1^*, q_2^*, \dots, q_n^*)$ 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 $\Phi(P^*) \leq \Phi(P)$, we have

$$\frac{\sum_{i=1}^n \mu_i q_i^*}{\sqrt{\sum_{i=1}^n \sigma_i^2 (q_i^*)^2}} \leq \Phi(P^*) \leq \Phi(P).$$

Thus, $(q_1^*, q_2^*, \dots, q_n^*)$ 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

$$\begin{aligned} \mathbb{E}[u^D] &\geq u_P \quad (\text{from the previous claim}) \\ &= f_P(q_1, q_2, \dots, q_n) \\ &\geq f_P(q_1^*, q_2^*, \dots, q_n^*) \quad (\text{since } (q_1^*, q_2^*, \dots, q_n^*) \text{ is feasible and } (q_1, q_2, \dots, q_n) \text{ is optimal}) \\ &= \sum_{i=1}^n \frac{g_i^D}{2} (P(1 - q_i^*) + (1 - P)(1 + q_i^*)) \\ &= \sum_{i=1}^n \frac{g_i^D}{2} (1 + q_i^*) - P \sum_{i=1}^n g_i^D q_i^* \\ &= \sum_{i=1}^n \frac{g_i^D}{2} (1 + q_i^*) - P^* \sum_{i=1}^n g_i^D q_i^* - (P - P^*) \sum_{i=1}^n g_i^D q_i^* \\ &= \sum_{i=1}^n \frac{g_i^D}{2} (1 + 2p_i^* - 1) - P^* \sum_{i=1}^n g_i^D (2p_i^* - 1) - (P - P^*) \sum_{i=1}^n g_i^D q_i^* \\ &= P^* \sum_{i=1}^n g_i^D (1 - p_i^*) + (1 - P^*) \sum_{i=1}^n g_i^D p_i^* - (P - P^*) \sum_{i=1}^n g_i^D q_i^* \\ &= \Pr[\text{chooser picks 1}] \sum_{i=1}^n g_i^D (1 - p_i^*) + \Pr[\text{chooser picks 2}] \sum_{i=1}^n g_i^D p_i^* - (P - P^*) \sum_{i=1}^n g_i^D q_i^* \\ &= u^* - (P - P^*) \sum_{i=1}^n g_i^D q_i^* \\ &\geq u^* - |(P - P^*)| \sum_{i=1}^n |g_i^D| |q_i^*| \\ &\geq u^* - \delta \sum_{i=1}^n |g_i^D| \quad (\text{since } P - \delta \leq P^* \text{ and each } |q_i| \leq 1) \\ &= u^* - \gamma \end{aligned}$$

as desired. □

B Proof of Theorem 5.1 (ii)

Let b be an upper bound on the support of \mathcal{D} . We may decompose \mathcal{D} into two distributions, \mathcal{D}_1 and \mathcal{D}_2 , such that every value in the support of \mathcal{D}_1 is less than or equal to every value in the support

of \mathcal{D}_2 , and sampling from \mathcal{D} is equivalent to drawing j uniformly from $\{1, 2\}$ and then sampling from \mathcal{D}_j . Let

$$\mu^+ := \mathbb{E}_{g \sim \mathcal{D}_2} [g],$$

$$\mu := \mathbb{E}_{g \sim \mathcal{D}} [g].$$

By the assumption that \mathcal{D} is supported on at least two different values, we must have that $\mu < \mu^+$. To prove that the divider is better off, it thus suffices to show that, according to the equilibrium strategies,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}_{g^D, g^C \sim \mathcal{D}^n} [u^D] \geq \frac{\mu^+}{2}, \quad (6)$$

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}_{g^D, g^C \sim \mathcal{D}^n} [u^C] \leq \frac{\mu}{2}. \quad (7)$$

(Truly, these will both be equalities, but that is not necessary for the present result.)

To prove (6), fix an arbitrary small $\varepsilon > 0$ and consider the following strategy for the divider. Supposing that values are sampled from \mathcal{D}_j for a random $j \in \{1, 2\}$, let E be the event that at least $\left\lceil \frac{\sqrt[3]{1-\varepsilon} \cdot n}{2} \right\rceil$ goods are drawn from \mathcal{D}_2 instead of \mathcal{D}_1 . It follows from the law of large numbers that, for sufficiently large n , $\Pr[E] \geq \sqrt[3]{1-\varepsilon}$. Assuming this happens, the divider places the top-valued $\left\lceil \frac{\sqrt[3]{1-\varepsilon} \cdot n}{2} \right\rceil$ entirely in pile 1 and all other goods entirely in pile 2. Let X^C be the chooser's value for pile 1 minus the chooser's value for pile 2. Observe that X^C is the sum of n independent random variables bounded between $-b$ and b , and

$$\mathbb{E}[X^C] = \left\lceil \frac{\sqrt[3]{1-\varepsilon} \cdot n}{2} \right\rceil \mu - \left(n - \left\lceil \frac{\sqrt[3]{1-\varepsilon} \cdot n}{2} \right\rceil \right) \mu \geq ((\sqrt[3]{1-\varepsilon} - 1)n - 2) \mu.$$

Therefore, by Hoeffding's inequality, the probability that the chooser picks pile 1 is

$$P = \Pr[X^C > 0] \leq \exp\left(-\frac{2((\sqrt[3]{1-\varepsilon} - 1)n - 2)\mu^2}{(2b)^2 n}\right).$$

Note that, as $n \rightarrow \infty$, $P \rightarrow 0$. Thus, let n be sufficiently large so that this probability is less than $(1 - \sqrt[3]{1-\varepsilon})$. Then, according to this divider strategy,

$$\begin{aligned} \mathbb{E}[u^D] &\geq \Pr[E] \cdot (1 - P) \cdot \mathbb{E}[u^D \mid E \text{ and the chooser picks pile 2}] \\ &\geq \sqrt[3]{1-\varepsilon} \cdot \sqrt[3]{1-\varepsilon} \cdot \left\lceil \frac{\sqrt[3]{1-\varepsilon} \cdot n}{2} \right\rceil \mu^+ \\ &\geq \frac{(1-\varepsilon)n\mu^+}{2}. \end{aligned}$$

This implies Equation (6).

We next claim that, for any $\varepsilon > 0$, the probability that it is optimal for the divider to pick a division p (with auxiliary q) such that

$$\left| \sum_{i=1}^n q_i \right| \geq \varepsilon n$$

vanishes as $n \rightarrow \infty$. Suppose p, q is such a division, and without loss of generality, assume the divider prefers pile 1 and the chooser is more likely to pick pile 2, which implies that, for all sufficiently large n ,

$$\sum_{i=1}^n q_i < -\varepsilon n.$$

This means there must be at least one good i^* for which $q_{i^*} < 0$, i.e., $p_{i^*} \leq \frac{1}{2}$. Let p', q' be the alternative division where $p'_{i^*} := p_{i^*} + \frac{1}{2}$, and all other goods are divided the same. Let $P, u^D, P', (u^D)'$ refer to the probability the chooser picks pile 1 and divider utility in the original division and new division, respectively. We may bound upper-bound the interim expected divider utility in the old division by assuming the divider receives his preferred pile 1, as

$$\mathbb{E}[u^D] \leq \sum_{i=1}^n p_i g_i^D.$$

On the other hand,

$$\mathbb{E}[(u^D)'] \geq (1 - P') \sum_{i=1}^n p'_i g_i^D = (1 - P') \left(\sum_{i=1}^n p_i g_i^D + \frac{g_{i^*}^D}{2} \right).$$

Thus, we may bound

$$\begin{aligned} \mathbb{E}[(u^D)'] - \mathbb{E}[u^D] &\geq (1 - P') \frac{g_{i^*}^D}{2} - P' \sum_{i=1}^n p_i g_i^D \\ &\geq \frac{g_{i^*}^D}{4} - P' \sum_{i=1}^n p_i g_i^D \quad (\text{since } P < \frac{1}{2} \implies P' < \frac{1}{2}) \\ &= \frac{g_{i^*}^D}{4} - \Pr \left[\sum_{i=1}^n q'_i g_i^C > 0 \right] \sum_{i=1}^n p_i g_i^D \quad (\text{since } P < \frac{1}{2} \implies P' < \frac{1}{2}) \\ &\geq \frac{g_{i^*}^D}{4} - \exp \left(\frac{-2 \left(\sum_{i=1}^n q'_i \mu \right)^2}{(2b)^2 n} \right) \sum_{i=1}^n p_i g_i^D \quad (\text{by Hoeffding's inequality}) \\ &\geq \frac{\min_i g_i^D}{4} - \exp \left(\frac{-2 \left(-n\varepsilon + \frac{1}{2} \mu \right)^2}{4b^2 n} \right) bn. \end{aligned}$$

By assumption, the expectation of the first term vanishes at most subexponentially as $n \rightarrow \infty$ while the second term vanishes exponentially. Hence, with high probability, the division p' was better than p , so our claim is proved.

Equation (8) now follows via a final application of Hoeffding's inequality to the chooser's value. In more detail, assume that

$$\left| \sum_{i=1}^n q_i \right| < \varepsilon n,$$

which we may assume happens with probability at least $1 - \varepsilon$ for sufficiently large n . Then the expected chooser value for pile 1 is

$$\sum_{i=1}^n p_i \mu = \sum_{i=1}^n \frac{(1 + q_i) \mu}{2} \leq \frac{\mu n}{2} + \left| \sum_{i=1}^n q_i \right| \frac{\mu}{2} < \frac{(1 + \varepsilon) n \mu}{2},$$

and the expected chooser value for pile 2 is likewise

$$\sum_{i=1}^n (1 - p_i)\mu = \sum_{i=1}^n \frac{(1 - q_i)\mu}{2} \leq \frac{\mu n}{2} + \left| \sum_{i=1}^n q_i \right| \frac{\mu}{2} < \frac{(1 + \varepsilon)n\mu}{2}.$$

It follows from a similar application of Hoeffding's inequality as before that, for sufficiently large n , with probability at least $1 - \varepsilon$, the chooser value for each pile will be at most $\frac{(1+2\varepsilon)n\mu}{2}$. By the union bound, with probability at least $1 - 3\varepsilon$, the chooser values for both piles will be at most $\frac{(1+2\varepsilon)n\mu}{2}$, while with probability at most 3ε , we may still bound the chooser value for the best pile by nb . Thus, for sufficiently large n , the total expected chooser value is at most

$$\mathbb{E}[u^C] \leq \frac{(1 - 3\varepsilon)(1 + 2\varepsilon)n\mu}{2} + 3\varepsilon nb.$$

This implies Equation (7), completing the proof of (ii). □