

# DIVIDING A CAKE FOR THE IRRATIONALLY ENTITLED

FLORIAN BRANDL\* AND ANDREW MACKENZIE<sup>† ‡</sup>

This draft: August 12, 2025

## Abstract

A perfectly divisible cake is to be divided among a group of agents. Each agent is entitled to a share between zero and one, and these entitlements are compatible in that they sum to one. The mediator does not know the preferences of the agents, but can query the agents to make cuts and appraise slices in order to learn. We prove that if one of the entitlements is irrational, then the mediator must use a protocol that involves an arbitrarily large number of queries in order to construct an allocation that respects the entitlements regardless of preferences.

**Keywords:** fair division, cake-cutting, query complexity, proportionality

**JEL Codes:** D63, D71

## 1 Introduction

### 1.1 Overview

Consider the problem of privatizing some jointly-owned property in accordance with well-defined ownership shares, provided that (i) dividing the property among the owners is more desirable than selling the property and splitting the proceeds, (ii) the owners are unable or unwilling to compensate one another using money, and (iii) whether or not a given allocation respects a given agent's share is a subjective matter that depends on that agent's preferences. For example, a plot of land is to be divided in accordance with a will among heirs who love the old family home, or a divorcing couple's cherished assets are to be divided in accordance with a prenuptial agreement. When the ownership shares are equal, this is the classic problem of fair division ([Steinhaus, 1948](#)). In this article, we investigate the communication costs associated with solving this problem through mediation when the agents' preferences are private information, building on the recent contributions of [Cseh and Fleiner \(2020\)](#). In order to distinguish these communication costs from incentive costs, we proceed under the first-best assumption that the agents will be honest regardless of their incentives, and even so we find that these costs may be unbounded: there may be no deadline by which a solution can be guaranteed.

\*Department of Economics, University of Bonn, Bonn, Germany. Email: [florian.brandl@uni-bonn.de](mailto:florian.brandl@uni-bonn.de)

<sup>†</sup>Department of Microeconomics and Public Economics, Maastricht University, Maastricht, the Netherlands. Email: [a.mackenzie@maastrichtuniversity.nl](mailto:a.mackenzie@maastrichtuniversity.nl)

<sup>‡</sup>Florian Brandl acknowledges support by the DFG under the Excellence Strategy EXC-2047.

The importance of considering communication costs when comparing institutions dates back at least to [Hayek \(1945\)](#), who argued that a crucial benefit of decentralized markets relative to central planning is their lower communication costs. For emphasis, these costs are distinct from other considerations related to private information, such as incentive compatibility, credibility, transparency, and privacy.<sup>1</sup> Roughly, the concern we investigate is not that strategically transmitted information may be inaccurate, or that the agents may know too little relative to the mediator throughout the process, or that sensitive information may be unnecessarily exposed; it is instead that mediation may be impractical simply because it may take too long.

There are several ways to measure communication costs. In the context of mechanism design, previous work has used *communication complexity*, or the worst-case volume of bits ([Yao, 1979](#); [Kushilevitz and Nisan, 1997](#)), to formalize that incentivizing honesty is costly and that costs can be mitigated by using multi-stage protocols ([Segal, 2007](#); [Fadel and Segal, 2009](#); [Segal, 2010](#)). In the context of fair division, costs are measured instead using *query complexity*, or the worst-case number of questions asked by the mediator. We show that mediation may be prohibitively expensive even without incentivizing honesty while using a multi-stage protocol, and we do so using query complexity. At a high level, the novelty of our contribution can be illustrated using the following two problems.

**PROBLEM 1:** A cake must be divided fairly across one million people. It is not enough for the mediator to ensure that each person measures the value of his serving to be at least one-millionth of the cake; the mediator must furthermore ensure that each person deems his serving to be at least as good as any other.

**PROBLEM 2:** A decedent leaves a plot of land to be divided across two descendants, a child and a grandchild, in accordance with an unusual will: the child’s share of the estate must be equal to the proportion of the grandchild’s share to the child’s share. The mediator need only ensure that each descendant measures the value of his parcel to be at least his share of the whole plot.

Which problem is harder, as measured by the deadline the mediator can guarantee to the agents, when the only activity that costs any time is communication by means of classic Robertson-Webb queries ([Robertson and Webb, 1998](#); [Woeginger and Sgall, 2007](#))? Perhaps surprisingly, the two-agent problem is harder. Indeed, for Problem 1, the mediator can guarantee that a solution is found in finite time ([Brams and Taylor, 1995](#)), and moreover can guarantee a deadline by which it is found ([Aziz and Mackenzie, 2016](#)). By contrast, for Problem 2, while the mediator can still guarantee that a solution is found in finite time ([Barbanel, 1995](#); [Shishido and Zeng, 1999](#); [Cseh and Fleiner, 2020](#)), we show that the mediator can never guarantee a deadline if there is an agent whose entitlement is an irrational number ([Corollary 1](#)), and this is the case for Problem 2.<sup>2</sup> Notably, this establishes the existence of a classic cake division problem that is finite but unbounded.

<sup>1</sup>Incentive compatibility is the central consideration in the large literature on mechanism design; see for example [Gibbard \(1973\)](#). For credibility, see [Akbarpour and Li \(2020\)](#). For transparency, see for example [Hakimov and Raghavan \(2023\)](#). For privacy in economics distinct from communication costs, see for example [Gradwohl and Smorodinsky \(2017\)](#), [Gradwohl \(2018\)](#), [Milgrom and Segal \(2020\)](#), [Eliaz, Eilat, and Mu \(2021\)](#), [Mackenzie and Zhou \(2022\)](#), and [Haupt and Hitzig \(2024\)](#).

<sup>2</sup>Indeed, the child’s share  $x \in [0, 1]$  must satisfy  $1 - x = x^2$ , or  $x = \frac{-1+\sqrt{5}}{2} \approx 61.8\%$ ; this is simply the inverse of the golden ratio.

Corollary 1, in turn, follows from our main result, which involves associating each entitlement profile with an index familiar to the literature: the least common denominator of the entitlements, defined as infinity when there is an irrational entitlement. We refer to this as the *clonage* to suggest the minimum number of clones required to form an economy with equal entitlements, provided we must replace each agent with some number of clones who equally split his entitlement.<sup>3</sup>

It is an old observation that we can solve a problem with finite clonage by solving the associated clone problem where each agent’s clones inherit his preferences, and that this in turn can be solved by a guaranteed deadline. Indeed, the idea can be illustrated simply using the Dubins-Spanier variant of the Banach-Knaster protocol (Steinhaus, 1948; Dubins and Spanier, 1961): given a knife that glides over the cake from left to right, (i) ask each of the  $n$  clones to make a cut such that he measures the value to the left of his cut to be  $\frac{1}{n}$ , (ii) give a clone with a left-most cut everything to the left of his cut, and (iii) iterate. We solve the original problem by giving each agent all slices assigned to his clones, and if the original entitlement profile has clonage  $n$ , then we require  $\sum_{i=1}^n i = \frac{n^2+n}{2}$  queries.

It is far less obvious, but true, that it is unnecessarily costly for the mediator to solve the associated clone problem. Indeed, costs can be reduced by using the more elaborate Even-Paz protocol (Even and Paz, 1984) in place of the simple Dubins-Spanier protocol, and in a suitable sense this is the optimal protocol for solving the clone problem (Edmonds and Pruhs, 2011). That said, Cseh and Fleiner (2020) show that costs can be reduced *further* by directly solving the original problem using a recursive version of the two-agent Cut Near-Halves protocol (Robertson and Webb, 1998), and that costs can be reduced *further still* by iteratively splitting the original problem into two sub-problems using the Cseh-Fleiner protocol (Cseh and Fleiner, 2020).

Even so, clonage remains relevant to the analysis. In particular, if there are  $n$  agents and the clonage level is  $c$ , then the Cseh-Fleiner protocol requires  $2(n-1)\lceil \log_2 c \rceil$  queries (Cseh and Fleiner, 2020). Curiously, then, even though the Cseh-Fleiner protocol does not solve the associated clone problem, it grows more costly with clonage. Why?

Our main result provides a general answer: even though it is unnecessarily costly to solve the associated clone problem, *clonage causes communication costs*. More precisely, for each clonage level  $c$ , and for each problem with  $n$  agents and a clonage above  $c$ , *no* protocol is guaranteed to finish within  $\lfloor \log_2 \log_2 2c^{\frac{1}{n-1}} \rfloor$  queries (Theorem 1). Our lower bound and the Cseh-Fleiner upper bound together establish that a sufficient increase in clonage alone can cause the mediator’s problem to become harder, though the gap between these bounds leaves open the question of whether clonage is responsible for *all* communication costs.

The explanation relies on an alternative interpretation of clonage: instead of viewing it as an index of how many *agent* clones we require, we view it as an index of how many *cake* clones we require. More precisely, the clonage  $c$  is the minimum number of miniature cake replicas required to solve the problem in zero queries, provided that each agent values each serving of a replica to be worth  $\frac{1}{c}$  of the associated serving in the original cake. This is an index of the difficulty the mediator faces at his initial information set, and we establish our main result by generalizing this difficulty index to other information sets.

<sup>3</sup>There is no word in English for *amount of clones*. We propose *clonage*, like coverage for an amount covered, leakage for an amount leaked, or mileage for an amount of miles, but unlike coinage for the creation of a new word.

axiom	lower bound	upper bound
no-envy	$\Omega(n^2)$ , <a href="#">Procaccia, 2009</a>	$\mathcal{O}(n^{n^{n^n}})$ , <a href="#">Aziz and Mackenzie, 2016</a>
proportionality	$\Omega(n \log_2 n)$ , <a href="#">Edmonds and Pruhs, 2011</a>	$\mathcal{O}(n \log_2 n)$ , <a href="#">Even and Paz, 1984</a>
$e$ -proportionality	$\Omega(\log_2 \log_2 C(e))$ , <a href="#">Corollary 2</a>	$\mathcal{O}(\log_2 C(e))$ , <a href="#">Cseh and Fleiner, 2020</a>

**Table 1:** *Query bounds for cake division.* We hold the cake and knives fixed. The number of agents is  $n$ , the clonage of entitlement profile  $e$  is  $C(e)$ , and for each problem the **cost** we bound is the worst-case number of queries across measure profiles. In the first two rows,  $n$  alone determines a unique problem. In the third row, an arbitrary number of agents  $n \geq 2$  is held constant, and the clonage determines a class of problems indexed by entitlement profiles; the lower bound is for the best-case **cost** across problems given the clonage and the upper bound is for the worst-case **cost** across problems given the clonage. See [Section 4.2](#) for the formal definitions.

## 1.2 Related literature

To facilitate this discussion, we first set the stage with some key definitions in the context of a classic concrete example.

**EXAMPLE 1:** The mediator must allocate the *interval cake*  $C = (0, 1]$  using the *standard knife*  $(\kappa_t)_{t \in [0,1]}$  defined as follows: for each  $t \in [0, 1]$ , we have  $\kappa_t = (0, t]$ . A *serving* is any finite union of intervals of the form  $(a, b]$ , and a *kitchen measure* is a probability measure that assigns a positive value to each nonempty serving. The set of *agents* is  $N = \{1, 2, \dots, n\}$ , the *entitlement profile* is  $e \in [0, 1]^N$  with  $\sum_{i \in N} e_i = 1$ , and an *allocation* is an assignment of servings to agents  $X = (X_i)_{i \in N}$  such that the servings are pairwise disjoint and collectively exhaustive. Initially, the mediator knows that the agents' preferences are represented by a kitchen measure profile  $\mu = (\mu_i)_{i \in N}$ , but does not know which. In the context of a kitchen measure profile  $\mu$ , an allocation  $X$  is (i) *proportional* if for each  $i \in N$  we have  $\mu_i(X_i) \geq \frac{1}{n}$ , (ii) *envy-free* if for each pair  $i, j \in N$  we have  $\mu_i(X_i) \geq \mu_i(X_j)$ , and (iii)  *$e$ -proportional* if for each  $i \in N$  we have  $\mu_i(X_i) \geq e_i$ .

For each entitlement profile  $e$ , we establish a lower bound on the worst-case number of *queries* that the mediator must ask in an adaptive multi-stage *protocol* that determines an  $e$ -proportional allocation regardless of the kitchen measure profile. Our lower bound is established with particularly powerful queries, and thus holds for the classic Robertson-Webb queries used throughout this discussion; see [Cseh and Fleiner \(2020\)](#) for a detailed comparison of other query models.

We begin by discussing query bounds for envy-free allocation, proportional allocation, and  $e$ -proportional allocation, in sequence; we then conclude with some remarks about contributions we make to the model. Recall that in asymptotic order notation,  $\mathcal{O}(f(x))$  is an upper bound and  $\Omega(f(x))$  is a lower bound; see [Section 4.2](#) for the familiar definitions.

**NO ENVY.** When there are two agents, a protocol used by Prometheus and Zeus in Hesiod's *Theogony*, which dates from the eighth or seventh century BC (see [Lowry, 1987](#); [Brams, Taylor, and Zwicker, 1995](#)), constructs an allocation that is not only proportional but moreover envy-free: one agent divides the cake into two parts he deems equal, then the other agent chooses. In our model, because we work with particularly powerful queries, this is a one-query protocol.

Calls for a three-agent envy-free protocol date to at least [Gamow and Stern \(1958\)](#), and shortly thereafter, such a protocol was independently discovered by John Selfridge

and John Conway; see [Brams and Taylor \(1995\)](#). The extension of the Selfridge-Conway protocol to even four agents remained a major open question for decades, until the arrival of the Brams-Taylor protocol for any number of agents ([Brams and Taylor, 1995](#)).

For each kitchen measure profile, the Brams-Taylor protocol uses a finite number of queries, and yet for each deadline the mediator might hope to guarantee, there is a kitchen measure profile that takes too long. To quote [Procaccia \(2020\)](#): “Consequently, as soon as Brams and Taylor solved the envy-free cake-cutting problem, they immediately launched a new problem to the top of fair division’s most wanted list: the existence of a bounded envy-free cake-cutting protocol.” In the meantime, it was established that any envy-free protocol requires at least  $\Omega(n^2)$  queries ([Procaccia, 2009](#)), which together with the Even-Paz protocol for proportional allocation ([Even and Paz, 1984](#)) established that envy-free allocation is more difficult than proportional allocation.

Finally, after enjoying decades as the leading candidate for a classic cake division problem that is finite but unbounded, the problem of envy-free allocation was ultimately classified as bounded with the arrival of the Aziz-Mackenzie protocol ([Aziz and Mackenzie, 2016](#)). We continue the conversation by classifying a *different* classic cake division problem as finite but unbounded, establishing that just as envy-free allocation is more difficult than proportional allocation, so too is  $e$ -proportional allocation more difficult than envy-free allocation.

**RATIONAL ENTITLEMENTS.** In the seminal contribution to fair division, [Steinhaus \(1948\)](#) introduced his own proportional protocol for three agents as well as the Banach-Knaster protocol for any number of agents. As discussed earlier, the latter has a particularly simple variant involving a moving knife ([Dubins and Spanier, 1961](#)), yielding the original upper bound for proportional allocation of  $\mathcal{O}(n^2)$ . This was then improved to  $\mathcal{O}(n \log_2 n)$  by the Even-Paz protocol ([Even and Paz, 1984](#)), which after some partial lower bound results ([Magdon-Ismail, Busch, and Krishnamoorthy, 2003](#); [Woeginger and Sgall, 2007](#)) was shown to be optimal ([Edmonds and Pruhs, 2011](#)).

The seminal contribution also remarked that the Banach-Knaster protocol can be adapted to rational shares and mused about protocol performance. To quote [Steinhaus \(1948\)](#): “The procedure described above applies also, under slight modifications, to the case of different (but rational) ideal shares. Interesting mathematical problems arise if we are to determine the minimal numbers of ‘cuts’ necessary for fair division.” Presumably, the slight modifications refer to solving the associated clone problem discussed earlier.

Decades later, two alternatives to cloning were proposed for two-agent problems: the Ramsey Partition protocol ([McAvaney, Robertson, and Webb, 1992](#)) and the Cut Near-Halves protocol ([Robertson and Webb, 1998](#)). The latter requires fewer queries than the former ([Robertson and Webb, 1998](#)), and a recursive generalization of the latter requires fewer queries than solving the associated clone problem with *its* optimal protocol ([Cseh and Fleiner, 2020](#)). Decades later still, the Cseh-Fleiner protocol was proposed as a further improvement, and it provides the current upper bound on the number of queries required for  $e$ -proportional allocation ([Cseh and Fleiner, 2020](#)). As discussed earlier, even though the Cseh-Fleiner protocol does not solve the associated clone problem, its upper bound grows with clonage.

Finally, [Cseh and Fleiner \(2020\)](#) provide a lower bound whose distinction from our lower bound is subtle but significant. In particular, Cseh-Fleiner provide a lower bound for each problem as a function of what we call its *fineness*: the smallest  $n \in \mathbb{N}$  such

that  $\frac{1}{n}$  is at most the smallest positive endowment.<sup>4</sup> This can be rephrased using clonage in the following sense: if for each clonage level we cherry-pick a problem that is as fine as possible, then the number of required queries is  $\Omega(\log_2 C(e))$ , where  $C(e)$  denotes the clonage of  $e$ . By contrast, we prove that if for each clonage level we cherry-pick a problem that is as *easy* as possible, then the number of required queries is  $\Omega(\log_2 \log_2 C(e))$ . These lower bounds are not logically related, and at first glance they may appear rather similar. That said, the strong conclusions that we draw in our overview require our lower bound. Indeed, for Problem 2, where the mediator must divide the decedent’s estate across his child and grandchild, the fineness level is three but the clonage level is infinity. See [Section 4](#) for further discussion.

**GENERAL ENTITLEMENTS.** The general problem of constructing an  $e$ -proportional was first considered by [Barbanel \(1995\)](#), and to date there are three finite protocols that solve this problem: the Barbanel protocol ([Barbanel, 1995](#)), the Shishido-Zeng protocol ([Shishido and Zeng, 1999](#)), and the *second* Cseh-Fleiner protocol ([Cseh and Fleiner, 2020](#)). All three protocols are unbounded, and we provide a simple explanation: *no* protocol that solves this problem is bounded.

**THE MODEL.** First, we provide a measure-theoretic generalization of [Example 1](#) that to our knowledge is novel. Since at least [Woodall \(1980\)](#), the standard approach has been to assert that the cake lives in Euclidean space alongside the Lebesgue measure, then use the Lebesgue measure to construct knives and articulate the associated preference requirements. By contrast, we directly specify the knives as the primitives from which the servings are constructed, then use the knives to articulate our preference requirements. This additional generality is largely an aesthetic choice, guided by the subjective principle is that unnecessary details are often distracting.

Second, we provide a preference foundation for kitchen measures. While there are existing foundations for the measures required for cake division ([Barbanel and Taylor, 1995](#)), we are able to use the additional structure that our algebra of servings inherits from our primitive knives to provide foundations that are closer to classic decision theory. In particular, ranking cake servings on the basis of preferences is mathematically equivalent to ranking events on the basis of beliefs, and it is already known that Savage’s classic foundation for representing beliefs with measures ([Savage, 1954](#)) extends from power sets to algebras ([Wakker, 1981](#); [Marinacci, 1993](#)). In order to obtain kitchen measure representation, we further impose that (i) a serving is null if and only if it is what we call a *sliver*, and (ii) preferences are continuous with respect to knives. The contribution is that with our additional assumptions, we are able to drop Savage’s technical *tightness* axiom, replacing it with two conditions that are somewhat easier to interpret ([Theorem 2](#)).

Finally, deviating from earlier literature, we include a formal definition of protocols; they are strategies for the mediator in a particular one-player extensive-form game against nature. This formality allows us to translate observations into the language of game

<sup>4</sup>Intriguingly, tiny entitlements have also proven troublesome for a mechanism design problem related to privatizing jointly-owned property in accordance with well-defined ownership shares. In particular, if two agents with private preference information are to allocate one indivisible object and monetary transfers are feasible, then the entitlement profile has the property that it is possible to implement efficient and voluntary allocations if and only if no agent has a tiny entitlement ([Cramton, Gibbons, and Klemperer, 1987](#)). (More generally, for more than two agents, no agent should have a huge entitlement.) Here, even zero counts as a problematic tiny entitlement because an agent who owns nothing can still pay; thus the classic [Myerson and Satterthwaite \(1983\)](#) impossibility theorem is a special case.

theory, including notably our observation of an interesting phenomenon in a variant two-player zero-sum game between the mediator and the *adversary*: when we restrict attention to pure strategies, the game has a value but no Nash equilibrium.

## 2 Problems

### 2.1 Kitchens

In our problem, a mediator must cut a cake into servings while constrained by the available technology for making cuts.

DEFINITION: A *kitchen* is a pair  $(C, \mathbb{K})$  that satisfies the following.

- $C$  is a nonempty set referred to as the *cake*.
- A *knife* is a strictly monotonic function  $\kappa : [0, 1] \rightarrow 2^C$ , where for each  $t \in [0, 1]$  we write  $\kappa_t$  instead of  $\kappa(t)$ . Equivalently,  $(\kappa_t)_{t \in [0, 1]} \in (2^C)^{[0, 1]}$  is a knife if for each pair  $t, t' \in [0, 1]$  such that  $t' > t$ , we have  $\kappa_t \subsetneq \kappa_{t'}$ .
- $\mathbb{K}$  is a nonempty set of knives referred to as the *knife block*. This represents the available technology for cutting the cake.
- A *simple slice* is a member of  $\bigcup_{\kappa \in \mathbb{K}} \bigcup_{t \in [0, 1]} \{\kappa_t, C \setminus \kappa_t\}$ , and a *slice* is any finite intersection of simple slices. We let  $\mathcal{S}^\cap \subseteq 2^C$  denote the collection of slices. By convention, we include  $C \in \mathcal{S}^\cap$  as the slice associated with the empty collection of simple slices.
- A *serving* is a finite union of slices. We let  $\mathcal{S} \subseteq 2^C$  denote the collection of servings. By convention, we include  $\emptyset \in \mathcal{S}$  as the serving associated with the empty collection of slices. Observe that each slice is a serving:  $\mathcal{S}^\cap \subseteq \mathcal{S}$ .

Observe that  $\mathcal{S}$ , together with its set-theoretic structure, is a Boolean algebra.<sup>5</sup>

We remark that the model has several alternative interpretations—for example, a slice might represent a parcel of land, a time slot on a schedule, or an uncertain event—but for brevity we stick to the cake interpretation.

Intuitively, if knife  $\kappa$  glides over the cake from time zero to time one, then it can make a cut by stopping at any time  $t \in [0, 1]$ , in which case the cake is split into  $\kappa_t$  and  $C \setminus \kappa_t$ . We will eventually require that agents declare  $\kappa_0$  to be equivalent to no cake while declaring  $\kappa_1$  to be equivalent to the entire cake, so it may already be useful to imagine that a knife glides across (essentially) the entire cake. For our problem, the construction of a particular serving may involve multiple knives, but we assume that the technology for constructing servings is constrained to the fixed inventory of knives given by the knife block, and as a result some subsets of the cake may not be feasible servings.

EXAMPLE 1: Let  $C = (0, 1]^2$ ; we refer to this as the *square cake*. Define the *standard vertical knife* (which moves horizontally),  $(\kappa_t)_{t \in [0, 1]}$ , and the *standard horizontal knife*

<sup>5</sup>More precisely,  $\mathcal{S}$  is the subalgebra of  $2^C$  generated by the simple slices. It is clear that the latter contains the former, and to see that the former contains the latter it suffices to write each member of the latter in disjunctive normal form.

(which moves vertically),  $(\kappa'_t)_{t \in [0,1]}$ , as follows: for each  $t \in [0, 1]$ ,  $\kappa_t \equiv \{(x, y) \in C | x \leq t\}$  and  $\kappa'_t \equiv \{(x, y) \in C | y \leq t\}$ . Finally, define the knife block  $\mathbb{K} \equiv \{(\kappa_t)_{t \in [0,1]}, (\kappa'_t)_{t \in [0,1]}\}$ . We refer to  $(C, \mathbb{K})$  as the *(standard) square cake kitchen*. In this case,  $\{(x, y) \in C | x = y\}$  is an example of a subset of the cake that is not a serving.

Before proceeding, we observe that we allow for algebras that are not  $\sigma$ -algebras. This point is not crucial to our main message, but it is relevant to some detailed remarks about our relationship with the literature, and this point can be illustrated using our previous example.

**EXAMPLE 2:** Let  $(C, \mathbb{K})$  be the square cake kitchen. We show that a countable family of servings may have a supremum in  $\mathcal{S}$  that is distinct from its supremum in  $2^C$ , and may also have no supremum in  $\mathcal{S}$ . Indeed, for each  $m \in \mathbb{N}$ ,<sup>6</sup> define  $S_m \equiv (0, 1 - \frac{1}{2^m}] \times (0, 1]$  and  $S'_m \equiv \cup_{m' \in \{1, 2, \dots, m\}} (\frac{1}{3^{m'}}, \frac{2}{3^{m'}}] \times [0, 1]$ . First, the family  $\{S_m\}_{m \in \mathbb{N}}$  has supremum  $(0, 1) \times (0, 1]$  in  $2^C$  and supremum  $(0, 1] \times (0, 1]$  in  $\mathcal{S}$ ; these are distinct and the former does not belong to  $\mathcal{S}$ . Second, the family  $\{S'_m\}_{m \in \mathbb{N}}$  has no supremum in  $\mathcal{S}$ ; thus  $\mathcal{S}$  is not a  $\sigma$ -algebra.

## 2.2 Kitchen measures

The cake will ultimately be consumed by a group of agents, each agent has preferences over servings, and we make assumptions about these preferences that imply they can be represented by exactly one measure with some additional structure. In this case, we say that the measure assigns values to servings, and we use these values to articulate whether or not a given allocation is satisfactory.

In our problem, the mediator does not know the agents' preferences. To make the mediator's problem easier—and thus to make our negative result stronger—we suppose that the mediator knows that for each agent, a slice is worth nothing if and only if it is what we call a *sliver*: a slice that some knife can cut no further, or a serving contained in such a slice, or a finite union of such servings.

**DEFINITION:** Fix a kitchen. We define slivers in three steps: (i) a slice  $S \in \mathcal{S}^\cap$  is a *slice-sliver* if there is  $\kappa \in \mathbb{K}$  such that for each  $t \in [0, 1]$ , we have  $S \cap \kappa_t = \emptyset$  or  $S \setminus \kappa_t = \emptyset$ , (ii) a *subslice-sliver* is a serving that is a subset of a slice-sliver, and (iii) a *sliver* is a finite union of subslice-slivers. We let  $\mathcal{S}^\circ \subseteq \mathcal{S}$  denote the collection of slivers; this is the ideal generated by the slice-slivers. By convention, we include  $\emptyset \in \mathcal{S}^\circ$  as the sliver associated with the empty collection of subslice-slivers.

In order to learn about preferences, the mediator may select an agent, a knife, a serving, and a target proportion, then request that the agent use the knife to cut a part from the serving such that the ratio of the part's value to the serving's value is the target proportion. To make the mediator's problem even easier—and thus to make our negative result even stronger—we suppose that the mediator knows that agents are always able to complete these requests. In particular, we assume that the mediator knows that each agent's preferences are represented by what we call a *kitchen measure*.

**DEFINITION:** Fix a kitchen. A *probability measure* is a function  $\mu_0 : \mathcal{S} \rightarrow [0, 1]$  such that (i)  $\mu_0(S) = 1$ , and (ii) for each pair  $A, B \in \mathcal{S}$  such that  $A \cap B = \emptyset$ , we have

<sup>6</sup>In this paper, we write  $\mathbb{N} = \{1, 2, \dots\}$  and  $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$ .

$\mu_0(A \cup B) = \mu(A) + \mu(B)$ . In this case, we say that  $\mu_0$  is moreover a *kitchen measure* if it satisfies the following conditions.

- *Null slivers.* For each  $S \in \mathcal{S}$ , we have  $\mu_0(S) = 0$  if and only if  $S \in \mathcal{S}^\circ$ .
- *Knife divisibility.* For each  $A \in \mathcal{S}$ , each  $p \in [0, 1]$ , and each  $\kappa \in \mathbb{K}$ , there is  $t \in [0, 1]$  such that  $\mu(A \cap \kappa_t) = p \cdot \mu(A)$ .

We let  $\mathbb{M} \subseteq [0, 1]^\mathcal{S}$  denote the collection of kitchen measures. We say that  $(C, \mathbb{K})$  is *tidy* if and only if  $\mathbb{M} \neq \emptyset$ .

To conclude this section, we first illustrate the definition of kitchen measure using the square cake, then discuss tidy kitchens, and finally discuss preference assumptions. We begin with the illustration.

**EXAMPLE 3:** Let  $(C, \mathbb{K})$  be the square cake kitchen. The restriction of the Lebesgue measure to  $\mathcal{S}$  is a kitchen measure, so the square cake kitchen is tidy. The only sliver is the empty serving, so in order for a probability measure to be a kitchen measure, null slivers requires that each nonempty serving is assigned a positive value.

To see that null slivers does not imply knife divisibility, let  $V$  denote the vertical line segment  $\{0.5\} \times [0, 1]$ . For emphasis,  $V$  is not a serving and thus not a sliver. For each  $S \in \mathcal{S}$ , (i) define  $\mu^*(S)$  to be the one-dimensional Lebesgue measure of  $S \cap V$ , (ii) define  $\mu^{**}(S)$  to be the two-dimensional Lebesgue measure of  $S$ , and (iii) define  $\mu(S) \equiv \frac{1}{2}\mu^*(S) + \frac{1}{2}\mu^{**}(S)$ . Then  $\mu$  is a probability measure that satisfies null slivers, yet it violates the knife divisibility: for the standard vertical knife  $\kappa$ , there is no  $t \in [0, 1]$  such that  $\mu(C \cap \kappa_t) = \frac{1}{2} \cdot \mu(C) = \frac{1}{2}$ .

To see that knife divisibility does not imply null slivers, define  $C^* \equiv (0, \frac{1}{2}] \times (0, \frac{1}{2}]$ , and for each  $S \in \mathcal{S}$ , define  $\mu(S)$  to be the two-dimensional Lebesgue measure of  $S \cap C^*$  multiplied by four. Then  $\mu$  is a probability measure that satisfies knife divisibility, yet it violates null slivers because  $(\frac{1}{2}, 1] \times [0, 1]$  is not a sliver but nevertheless is assigned zero value.

As the previous example illustrates, the square kitchen is tidy. As the next example illustrates, not all kitchens are tidy.

**EXAMPLE 4:** Let  $C = [0, 1]^\mathbb{R}$ , and for each dimension  $d \in \mathbb{R}$ , let  $\kappa^d$  be the knife such that for each  $t \in [0, 1]$  we have  $\kappa_t^d = \{x \in C \mid x_d \leq t\}$ . We claim that  $(C, \mathbb{K})$  is not tidy. Indeed, for each  $d \in \mathbb{R}$ , define

$$A_d \equiv \{x \in C \mid x_d > \frac{1}{2} \text{ and for each } d' \in \mathbb{R} \setminus \{d\} \text{ we have } x_{d'} \leq \frac{1}{2}\}.$$

Then  $\{A_d\}_{d \in \mathbb{R}}$  is a continuum of pairwise disjoint servings that are not slivers, so there cannot be a kitchen measure. Indeed, assume by way of contradiction that  $\mu \in \mathbb{M}$ , and for each  $m \in \mathbb{N}$ , define  $\mathcal{D}_m \equiv \{d \in \mathbb{R} \mid \mu(A_d) \in (\frac{1}{m+1}, \frac{1}{m}]\}$ . Then (i)  $\cup_{m \in \mathbb{N}} \mathcal{D}_m = \mathbb{R}$ , so  $|\cup_{m \in \mathbb{N}} \mathcal{D}_m| = |\mathbb{R}|$ , and (ii) for each  $m \in \mathbb{N}$  we have  $|\mathcal{D}_m| \leq m$ , so  $|\cup_{m \in \mathbb{N}} \mathcal{D}_m| \leq |\mathbb{N}|$ . But then  $|\mathbb{R}| \leq |\mathbb{N}|$ , contradicting  $|\mathbb{N}| < |\mathbb{R}|$ .

In order for a kitchen to be tidy, the quotient algebra—obtained by declaring two servings to be equivalent if their symmetric difference is a sliver, gathering the equivalence

classes, and deriving algebraic operations from the original algebra in the natural way—must carry a strictly positive measure and therefore satisfy Kelley’s criterion (Kelley, 1959). The cake division literature has largely focused on tidy kitchens obtained like the square cake kitchen: begin with a probability space equipped with a countably additive and atomless reference measure—usually Euclidean space equipped with the Lebesgue measure—and then require each knife’s nested family of servings to be strictly monotonic in the reference measure. We remark that the inquiry into when certain algebras are compatible with certain measures traces back to at least 1937, when John von Neumann raised Problem 163 of the Scottish Book (Mauldin, 1981).

We conclude this section by discussing preference assumptions. To reiterate, we make assumptions about each agent’s preferences that imply they can be represented by exactly one kitchen measure. What, exactly, are these assumptions? The short answer is that we are assuming (i) the classic qualitative probability axioms (Bernstein, 1917; de Finetti, 1937; Koopman, 1940), (ii) one of the two technical axioms of Savage (1954), (iii) the null servings are the slivers, and (iv) each knife is continuous in a suitable sense. Notably, these assumptions together imply the second technical axiom of Savage (1954). These claims require further elaboration, but as they are not crucial to our main result about communication costs, we postpone the details until Section 5.1.

## 2.3 Settings and entitlements

The cake will ultimately be consumed by a group of agents, and the mediator must partition the cake into servings and then match these servings to the agents so that the resulting allocation is, in a sense to be made precise shortly, satisfactory.

DEFINITION: A *setting* is a tuple  $(C, \mathbb{K}, N)$  that satisfies the following.

- $(C, \mathbb{K})$  is a tidy kitchen.
- $N$  is a finite and nonempty set of *agents*, with  $n \equiv |N|$ .
- $\mathcal{X} \subseteq \mathcal{S}^N$  is the set of *allocations*: for each  $X \in \mathcal{S}^N$ , we have  $X \in \mathcal{X}$  if and only if (i) for each pair  $i, j \in N$ , we have  $X_i \cap X_j = \emptyset$ , and (ii)  $\cup_{i \in N} X_i = C$ .

All of our results involve fixing an arbitrary setting.

In our problem, each agent is equipped with both (i) an entitlement to a share of the cake between zero and one, and (ii) a kitchen measure representing his preferences. Moreover, the agents’ entitlements are compatible in that they sum to one. Whether or not a given allocation is satisfactory depends on both the entitlements and the kitchen measures: each agent should receive a serving that is worth at least his entitlement according to his measure.

DEFINITION: Fix a setting.

- $E \equiv \{e \in [0, 1]^N \mid \sum e_i = 1\}$  is the set of *entitlement profile*.
- $\mathbb{D} \equiv \mathbb{M}^N$  is the *unrestricted domain* (of all kitchen measure profiles).

For each  $e \in E$  and each  $\mu \in \mathbb{D}$ , an allocation  $X \in \mathcal{X}$  is  $(e|\mu)$ -proportional if for each  $i \in N$ , we have  $\mu_i(X_i) \geq e_i$ . When omitting  $\mu$  creates no confusion, we use the simpler term  $e$ -proportional.

In the special case that each agent has entitlement  $\frac{1}{n}$ ,  $e$ -proportionality coincides with the classic notion of proportionality (Steinhaus, 1948). As discussed earlier, the general existence of  $e$ -proportional allocations has been shown by construction (Barbanel, 1995; Shishido and Zeng, 1999; Cseh and Fleiner, 2020), and there are particularly simple proofs of existence for two special cases: (i) each agent’s measure is countably additive (Liapounoff, 1940; Dubins and Spanier, 1961), and (ii) each entitlement is a rational number (Steinhaus, 1948; Dubins and Spanier, 1961).

We consider a mediator who knows the entitlement profile but not the kitchen measure profile, and who is tasked with constructing an  $e$ -proportional allocation *regardless of the kitchen measure profile*. To do so, the mediator may communicate with the agents in order to learn more about their preferences, and the allocation that the mediator constructs may be contingent on the information that he receives. That said, communication is costly, and the mediator’s problem is to keep these costs down.

## 3 Solutions

### 3.1 Queries

In order to learn about the agents’ preferences, the mediator selects a *query* and receives a *response*. Since the seminal contribution of Robertson and Webb (1998), these terms have been formalized in several different ways (Woeginger and Sgall, 2007; Edmonds and Pruhs, 2011; Cseh and Fleiner, 2020).

Our choice of definition is based on two considerations. First, as an aesthetic choice, we require that whenever a cut is made, the mediator does not lose access to old queries or gain access to new queries: the mediator adaptively selects from the fixed set of queries. Second, because our main result loosely states that many queries are required to solve our problem, we should allow for ‘powerful’ queries in order to strengthen our main result.

**DEFINITION:** Fix a setting. A *query* is a tuple  $\mathbf{q} = (i, \kappa, S, p)$  specifying (i) a *cutter*  $i \in N$ , (ii) a knife  $\kappa \in \mathbb{K}$ , (iii) a serving  $S \in \mathcal{S}$ , and (iv) a target proportion  $p \in [0, 1]$ . The interpretation is as follows:

- First, the mediator uses the knives to construct  $S$ .
- Second, the mediator asks  $i$  to use both his private measure and the knife  $\kappa$  to cut  $S$  in order to construct  $S' \subseteq S$  with the property that, according to his private measure, the proportion of the value of  $S'$  to the value of  $S$  is precisely the target proportion  $p$ . To be explicit, if this private measure is  $\mu_i$ , then the constructed  $S'$  should satisfy  $\mu_i(S') = p \cdot \mu_i(S)$ . Furthermore, the mediator provides detailed instructions about how  $S'$  should be constructed; see the next section.
- Third, the mediator asks every agent to use his private measure to appraise both the value of  $S$  and the value of  $S'$ .

We let  $\mathcal{Q}$  denote the set of queries.

This query definition is our adaptation of the ‘stronger query definition’ of [Cseh and Fleiner \(2020\)](#) to our model.

### 3.2 The cutter’s serving

For each query, the cutter is able to construct a serving that satisfies the mediator’s requirements because his preferences are represented by a kitchen measure. Moreover, if there are several such servings, then it does not matter which the cutter selects: in all cases the mediator receives the same information because he knows which servings are null. It is therefore without loss of generality to assert that the cutter is given specific instructions for how to break ties, and doing so simplifies the presentation.

**DEFINITION:** Fix a setting. For each query  $\mathbf{q} = (i, \kappa, S, p)$  and each  $\mu_i \in \mathbb{M}$ , we define the  $(\mathbf{q}|\mu_i)$ -threshold by  $\tau(\mathbf{q}|\mu_i) \equiv \inf\{t \in [0, 1] \mid \mu_i(S \cap \kappa_t) \geq p \cdot \mu_i(S)\}$  and we define the  $(\mathbf{q}|\mu_i)$ -serving by  $S_{\mathbf{q}|\mu_i} \equiv S \cap \kappa_{\tau(\mathbf{q}|\mu_i)}$ . The interpretation is that when the mediator selects query  $\mathbf{q}$  and the cutter  $i$  has preferences represented by  $\mu_i$ , the output serving is  $S_{\mathbf{q}|\mu_i}$ .

The intended interpretation is that in order to construct  $S_{\mathbf{q}|\mu_i}$ , the cutter  $i$  glides  $\kappa$  over  $S$ , starting at time zero, until the first moment that  $S \cap \kappa_t$  is worth at least the target value, at which point he cuts. By the following lemma, the intended interpretation is valid.

**LEMMA 1:** Fix a setting. For each query  $\mathbf{q} = (i, S, \kappa, p)$  and each  $\mu \in \mathbb{D}$ , the constructed serving satisfies the mediator’s request:  $\mu_i(S_{\mathbf{q}|\mu_i}) = p \cdot \mu_i(S)$ .

**PROOF:** Assume the hypotheses. Since  $\mu_i$  is a kitchen measure, there is  $t \in [0, 1]$  such that  $\mu_i(S \cap \kappa_t) = p \cdot \mu_i(S)$ . Then  $\tau(\mathbf{q}|\mu_i) \leq t$ , so  $S_{\mathbf{q}|\mu_i} = S \cap \kappa_{\tau(\mathbf{q}|\mu_i)} \subseteq S \cap \kappa_t$ , so  $\mu_i(S_{\mathbf{q}|\mu_i}) \leq \mu_i(S \cap \kappa_t) = p \cdot \mu_i(S)$ .

Assume, by way of contradiction, that  $\mu_i(S_{\mathbf{q}|\mu_i}) < p \cdot \mu_i(S)$ . Then there is  $p' \in [0, 1]$  such that  $\mu_i(S_{\mathbf{q}|\mu_i}) < p' \cdot \mu_i(S) < p \cdot \mu_i(S)$ , so as  $\mu_i \in \mathbb{M}$  there is  $t' \in [0, 1]$  such that  $\mu_i(S \cap \kappa_{t'}) = p' \cdot \mu_i(S) > \mu_i(S_{\mathbf{q}|\mu_i})$ , so  $t' > \tau(\mathbf{q}|\mu_i)$ . But then by definition of  $\tau(\mathbf{q}|\mu_i)$  we have  $\mu_i(S \cap \kappa_{t'}) \geq p \cdot \mu_i(S)$ , contradicting  $\mu_i(S \cap \kappa_{t'}) = p' \cdot \mu_i(S) < p \cdot \mu_i(S)$ . Altogether, then,  $\mu_i(S_{\mathbf{q}|\mu_i}) = p \cdot \mu_i(S)$ , as desired. ■

### 3.3 Records, restrictions, and responses

Initially, the mediator knows only that the measure profile belongs to the unrestricted domain  $\mathbb{D}$ , but after selecting a query and receiving the agents’ response, the mediator learns that the measure profile belongs to a restricted domain  $\mathbb{D}' \subseteq \mathbb{D}$ . More generally, from a list of queries and their responses, the mediator learns a restricted domain. We begin with the general definitions and then define the response to an individual query as a special case.

**DEFINITION:** Fix a setting. For each collection of servings  $\mathcal{S}' \subseteq \mathcal{S}$ , an  $\mathcal{S}'$ -appraisal is a function  $\alpha_0 : \mathcal{S}' \rightarrow [0, 1]$ . A (finite) record is a pair  $(\mathcal{S}', \alpha')$  such that (i)  $\mathcal{S}' \subseteq \mathcal{S}$  with  $|\mathcal{S}'| \in \mathbb{N}$ , and (ii)  $\alpha' = (\alpha'_i)_{i \in N}$  is a profile of  $\mathcal{S}'$ -appraisals. In this case, we define the associated restricted domain to be the collection of kitchen measure profiles that extend

the appraisal profile,  $\mathbb{D}(\mathcal{S}', \alpha') \equiv \{\mu \in \mathbb{D} \mid \text{for each } i \in N, \alpha'_i = \mu_i \upharpoonright_{\mathcal{S}'}\}$ , and we say that  $(\mathcal{S}', \alpha')$  is *valid* if  $\mathbb{D}(\mathcal{S}', \alpha') \neq \emptyset$ .<sup>7</sup>

DEFINITION: Fix a setting. Given a pair of records  $(\mathcal{S}', \alpha')$  and  $(\mathcal{S}'', \alpha'')$ , we say that  $(\mathcal{S}'', \alpha'')$  is *at least as informative as*  $(\mathcal{S}', \alpha')$  if  $\mathbb{D}(\mathcal{S}'', \alpha'') \subseteq \mathbb{D}(\mathcal{S}', \alpha')$ .

Initially, the mediator does not know how any agent appraises any non-null serving, but at any later stage the mediator will have some query responses that allow him to fill in some of the blanks. The resulting record is a collection of servings together with their appraisal from all agents, and this record in turn functions as a domain restriction: the mediator is sure that the measure profile extends the appraisal profile. A special case is the record consisting of a response to a single query.

DEFINITION: Fix a setting. For each query  $\mathbf{q} = (i, \kappa, S, p)$  and each  $\mu \in \mathbb{D}$ , the  $(\mathbf{q} \mid \mu)$ -response,  $r(\mathbf{q} \mid \mu)$ , is the record  $(\mathcal{S}', \alpha')$  such that (i)  $\mathcal{S}' = \{S, S_{\mathbf{q} \mid \mu_i}\}$ , and (ii) for each  $j \in N$ ,  $\alpha'_j = \mu_j \upharpoonright_{\mathcal{S}'}$ . In this case, the  $(\mathbf{q} \mid \mu)$ -domain,  $\mathbb{D}(\mathbf{q} \mid \mu)$ , is  $\mathbb{D}(r(\mathbf{q} \mid \mu))$ .

This definition fits our earlier discussion of queries. Indeed, suppose the mediator selects  $\mathbf{q}$  and the measure profile is  $\mu$ . Then first, the mediator constructs  $S$ ; second, the cutter constructs  $S_{\mathbf{q} \mid \mu_i}$ ; and third, every agent appraises both  $S$  and  $S_{\mathbf{q} \mid \mu_i}$ . As a result, the mediator learns that the measure profile belongs to  $\mathbb{D}(\mathbf{q} \mid \mu) \subseteq \mathbb{D}$ .

### 3.4 Protocols

The mediator must communicate with the agents until he can identify an allocation that is surely satisfactory, in the sense that it respects the entitlements, given what he has learned. Communication is costly, however, and thus the mediator seeks to select a (communication) protocol that keeps costs down, subject to the constraint that it always identifies a satisfactory allocation.

We formalize a protocol as a pure strategy for the mediator in the following game against nature: nature opens with an unobserved selection of the measure profile, and thereafter the mediator iteratively selects queries and receives responses until he ends the game by selecting an allocation. We begin by formalizing our game against nature, minimizing the notation that we introduce in anticipation of our proof's requirements.

DEFINITION: Fix a setting. For each  $e \in E$ , the *division game given e* is the one-player game against nature defined as follows.

- The only player with payoffs is the *mediator*, and *nature* also plays. There are also histories partially ordered by precedence, active players and actions associated with non-terminal histories, information sets, and payoffs. The precedence order is such that each history and action at that history together determine an *immediate successor*, or an earliest history that follows. We define these remaining ingredients below, defining precedence entirely using immediate successors.
- The initial history is non-terminal, its active player is nature, and its set of available actions is  $\mathbb{M}^N$ . Nature's choice of measure profile is not observed by the mediator.

<sup>7</sup>We are mostly interested in valid records, but for some arguments it is convenient to work with records that are not valid because they assign zero to servings that are not slivers.

After each choice of nature, the immediate successor is non-terminal.

- At each non-terminal history other than the initial history, the active player is the mediator and the set of available actions is  $\mathcal{Q} \cup \mathcal{X}$ . We denote the set of histories where the mediator is the active player by  $H$ . If the mediator selects a query, then the immediate successor is non-terminal. If the mediator selects an allocation, then the immediate successor is terminal.
- Each non-initial history  $h$  is identified with the finite sequence of actions that leads to  $h$  from the initial history: there is a (*query*) *count*  $c \in \mathbb{N}_0$  such that  $h = (\mu, \mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_c)$ . In this case, the *chronicle at*  $h$  is the list of queries and their responses,  $(\mathbf{q}_1, \mathbf{r}(\mathbf{q}_1|\mu), \mathbf{q}_2, \mathbf{r}(\mathbf{q}_2|\mu), \dots, \mathbf{q}_c, \mathbf{r}(\mathbf{q}_c|\mu))$ . Moreover, the *h-domain* is  $\mathbb{D}(h) \equiv \cap_{t \in \{1, 2, \dots, c\}} \mathbb{D}(\mathbf{r}(\mathbf{q}_t|\mu))$ , where by convention  $c = 0$  implies  $\mathbb{D}(h) = \mathbb{D}$ ; this is nonempty as  $\mu \in \mathbb{D}(h)$ . Finally, two histories share an *information set* if and only if (i) the mediator is the active player at both histories, and (ii) the two histories have the same chronicle. Observe that if  $h$  and  $h'$  share an information set, then  $\mathbb{D}(h) = \mathbb{D}(h')$ .
- A play is a maximal chain of histories. If a play ends at a terminal history with a query count of  $c$ , if the measure profile selected by nature is  $\mu$ , and if the allocation selected by the mediator is  $(e|\mu)$ -proportional, then the mediator's *payoff* is  $-c$ ; otherwise it is  $-\infty$ .

In this case, when we delete the payoffs, the result is the *division game form*. Note that our main result about communication costs does not directly involve these payoffs, but rather the related costs that we introduce shortly; thus at this point the payoffs can be taken as suggestive.

A protocol is a plan for the mediator in the division game form that is contingent only upon the available information.

DEFINITION: Fix a setting. A *protocol* is a pure strategy for the mediator in the division game form. Equivalently, a protocol is a function  $\pi : H \rightarrow \mathcal{Q} \cup \mathcal{X}$  such that for each pair  $h, h' \in H$ , if  $h$  and  $h'$  share an information set, then  $\pi(h) = \pi(h')$ . In this case, for each  $\mu \in \mathbb{D}$ , we define  $\text{cost}(\pi|\mu) \in \mathbb{N}_0 \cup \{\infty\}$  and  $\text{out}(\pi|\mu) \in \mathcal{X} \cup \{\text{error}\}$  as follows.

- If the play determined by  $\mu$  and  $\pi$  is infinite, then we define  $\text{cost}(\pi|\mu) \equiv \infty$  and  $\text{out}(\pi|\mu) \equiv \text{error}$ .
- If the play determined by  $\mu$  and  $\pi$  is finite, then let  $c$  denote the query count of its terminal history and let  $X$  denote its final action. In this case, we define  $\text{cost}(\pi|\mu) \equiv c$  and  $\text{out}(\pi|\mu) \equiv X$ .

We let  $\Pi$  denote the set of protocols. Observe that for each  $e \in E$ , the mediator's payoff in the division game given  $e$  for the play determined by  $\mu$  and  $\pi$  is (i)  $-\infty$ , if either  $\text{out}(\pi|\mu) = \text{error}$  or  $\text{out}(\pi|\mu)$  is not  $(e|\mu)$ -proportional, and (ii)  $-\text{cost}(\pi|\mu)$ , otherwise.

We assess each protocol in the context of a given entitlement profile  $e$ , but instead of doing so using payoffs directly, we separately assess (i) whether or not the protocol always constructs  $e$ -proportional allocations, and (ii) the worst-case communication cost measured in queries. We begin with the former.

DEFINITION: Fix a setting. For each  $e \in E$ , a protocol is *e-proportional* if for each  $\mu \in \mathbb{M}^N$ ,  $\text{out}(\pi|\mu) \in \mathcal{X}$  and  $\text{out}(\pi|\mu)$  is  $(e|\mu)$ -proportional. We let  $\Pi_e \subseteq \Pi$  denote the set of *e-proportional* protocols.

The set of *e-proportional* protocols is known to be nonempty (Barbanel, 1995; Shishido and Zeng, 1999; Cseh and Fleiner, 2020), and we are interested in comparing these protocols according to the following performance metric.

DEFINITION: Fix a setting and a protocol. The *worst-case cost* of  $\pi$  is

$$\text{cost}(\pi) \equiv \sup_{\mu \in \mathbb{D}} \text{cost}(\pi|\mu).$$

If  $\text{cost}(\pi) = \infty$ , then we say that  $\pi$  is *unbounded*; otherwise we say that  $\pi$  is *bounded*.

Altogether, then, the mediator's problem is given by a setting and an associated entitlement profile  $e$ , and the objective is to minimize the worst-case cost among the *e-proportional* protocols.

DEFINITION: Fix a setting. For each  $e \in E$ , the *optimal worst-case cost of an e-proportional protocol* is

$$\text{cost}(e) \equiv \min\{c \in \mathbb{N}_0 \cup \{\infty\} \mid \text{there is } \pi \in \Pi_e \text{ such that } \text{cost}(\pi) = c\}.$$

Because the set of *e-proportional* protocols is nonempty (Barbanel, 1995; Shishido and Zeng, 1999; Cseh and Fleiner, 2020), this is well-defined.

## 4 Communication costs

### 4.1 The mediator's problem

In their important contribution, Cseh and Fleiner (2020) associate each problem that the mediator might face with both a lower bound and an upper bound on its cost. These bounds and our analysis together involve three indices for entitlement profiles: clonage, precision, and fineness. The most important index is clonage, which we ultimately use to present both our lower bound and the Cseh-Fleiner upper bound. Our lower bound analysis involves precision, while the Cseh-Fleiner lower bound involves fineness, and crucially clonage can only be bounded as a function of the former.

If we restrict attention to entitlement profiles with rational and positive entitlements, then the three indices admit simple verbal descriptions.

- The *clonage* is the minimum number of clones required to form an equal-entitlement economy, provided we must replace each agent  $i$  with  $n_i \in \mathbb{N}$  clones who equally split  $e_i$ . This is the least common multiple of the denominators of the entitlements' reduced fractions.
- The *precision* is the maximum of the denominators of the entitlements' reduced fractions.
- The *fineness* is the smallest denominator of a positive fraction that is at most the smallest positive entitlement.

The formal definitions, which apply to all entitlement profiles, are as follows.

DEFINITION: Fix a setting. For each  $m \in \mathbb{N}$ , define the set of  $m$ -ruler markings by  $\mathcal{M}_m \equiv \{\frac{a}{m} | a \in \{0, 1, \dots, m\}\}$  and the set of *positive*  $m$ -ruler markings by  $\mathcal{M}_m^+ \equiv \mathcal{M}_m \setminus \{0\}$ . For each  $e \in E$ , we define the following.

- The *clonage* of  $e$  is  $C(e) \equiv \min(\{m \in \mathbb{N} | \text{for each } i \in N, e_i \in \mathcal{M}_m\} \cup \{\infty\})$ .
- For each  $i \in N$ ,  $P_i(e) \equiv \min(\{m \in \mathbb{N} | e_i \in \mathcal{M}_m\} \cup \{\infty\})$ . The *precision* of  $e$  is  $P(e) \equiv \max_{i \in N} P_i(e)$ .
- For each  $i \in N$ ,  $F_i(e) \equiv \min\{m \in \mathbb{N} | e_i \in \{0\} \cup [\min \mathcal{M}_m^+, 1]\}$ . The *fineness* of  $e$  is  $F(e) \equiv \max_{i \in N} F_i(e)$ .

When  $C$ ,  $P$ , and  $F$  are viewed as functions from  $E$  to  $\mathbb{N} \cup \{\infty\}$ , we refer to them as the *clonage index*, the *precision index*, and the *fineness index*, respectively.

The Cseh-Fleiner upper bound involves clonage.

THEOREM CF1:<sup>8</sup> Fix a setting. For each  $e \in E$  and each clonage level  $c \in \mathbb{N}$ ,  $C(e) = c$  implies  $\text{cost}(e) \leq 2(n-1) \lceil \log_2 c \rceil$ .

By contrast, at the entitlement profile level, the Cseh-Fleiner lower bound involves fineness.

THEOREM CF2:<sup>9</sup> Fix a setting. For each  $e \in E$  and each fineness level  $f \in \mathbb{N}$ ,  $F(e) > f$  implies  $\text{cost}(e) > (n-1) \log_3 f$ .

We sketch the intuition for this lower bound, which remarkably applies even when (i) there are two agents, (ii) we know one agent's measure, and (iii) after each query we learn the other agent's appraisal of the new partition generated by the old partition and the query. Indeed, suppose we know  $\mu_1$ , and suppose the entitlement profile has fineness greater than  $f$  because  $\frac{1}{f} > e_2 > 0$ . Any query and response together divide each cell of the old partition into at most three new sub-cells, and Cseh and Fleiner (2020) prove that it is always possible for 2 to respond that all of each old cell's value—except, perhaps, a negligible amount<sup>10</sup>—resides entirely in its new sub-cells that 1 values most. In this case, after  $q$  queries, a non-negligible cell for 2 is worth at least  $\frac{1}{3^q}$  to 1, we must give such a cell  $S$  to 2, and in order for 1 to agree we must have  $\frac{1}{3^q} \leq \mu_1(S) \leq 1 - e_1 = e_2 < \frac{1}{f}$ ; thus the number of queries  $q$  must be great enough given the fineness  $f$ .

Our first proposition provides a lower bound at the entitlement profile level using precision instead of fineness.

PROPOSITION 1: Fix a setting. For each  $e \in E$  and each precision level  $p \in \mathbb{N}$ ,  $P(e) > p$  implies  $\text{cost}(e) > \lceil \log_2 \log_2 2p \rceil$ .

<sup>8</sup>This is Theorem 5.5 of Cseh and Fleiner (2020).

<sup>9</sup>This is a variant of Theorem 6.3 in Cseh and Fleiner (2020) involving strict inequalities instead of weak inequalities; a minor modification of the original proof suffices for this version.

<sup>10</sup>This modification is needed in our model because the mediator knows the set of null servings precisely.

The proof spans [Appendix 1](#), [Appendix 2](#), and [Appendix 3](#). At a high level, the first appendix establishes our main two-agent lemma for queries that select a cell of the current partition ([Lemma 3](#)), the second appendix establishes our generalization of [Lemma 3](#) to all queries ([Lemma 4](#)), and the third appendix applies [Lemma 4](#) to establish the lower bound for any number of agents ([Proposition 1](#)). The key idea is that by cloning the cake instead of the agents, we can index the deficiency of a collection of responses. We discuss this proof in more detail in [Section 4.3](#).

The original presentation of the Cseh-Fleiner lower bound involves clonage, and we include such a statement as part of [Corollary CF](#) in [Section 4.2](#). Crucially, however, that statement is at the setting level. At the entitlement profile level, a lower bound involving fineness cannot be translated to a lower bound involving clonage, yet a lower bound involving precision can, and this is a consequence of the relationships between the three indices described by the following proposition.

**PROPOSITION 2:** The three indices can be compared as follows.

- For each setting and each  $e \in E$ , we have  $C(e) \geq P(e) \geq F(e)$ .
- There are a setting and an  $e \in E$  such that  $C(e) > P(e)$ .
- For each setting with  $n \geq 2$  and each  $e \in E$ ,  $C(e) = \infty$  implies  $P(e) = \infty$  and  $C(e) \in \mathbb{N}$  implies  $P(e)^{n-1} \geq C(e)$ .
- For each setting with  $n \geq 2$ , there is no  $f : \mathbb{N} \rightarrow \mathbb{N}$  such that for each  $e \in E$ ,  $P(e) \in \mathbb{N}$  implies  $f(F(e)) \geq P(e)$ .

**PROOF:** We begin with the first two items. The first item is immediate from definitions; we omit the proof. For the second item, an example is any setting with  $N = \{1, 2, 3, 4\}$ ,  $e_1 = \frac{1}{2}$ ,  $e_2 = \frac{1}{3}$ ,  $e_3 = \frac{1}{10}$ , and  $e_4 = \frac{1}{15}$ ; in this case  $C(e) = 30$  and  $P(e) = 15$ .

For the third item, fix a setting with  $n \geq 2$  and let  $e \in E$ . If  $C(e) = \infty$ , then there is an irrational entitlement and thus  $P(e) = \infty$ , as desired. If  $C(e) \in \mathbb{N}$ , then define  $c \equiv C(e)$  and define  $p \equiv P(e)$ . Moreover, for each  $i \in N$ , define  $b_i \equiv P_i(e)$ . Re-index the agents using  $N = \{1, 2, \dots, n\}$  such that  $b_1 \geq b_2 \geq \dots \geq b_n$ , then define  $b \equiv b_1 \cdot b_2 \cdot \dots \cdot b_{n-1}$ . For each  $i \in N \setminus \{n\}$ , there is  $a_i \in \{0, 1, \dots, b_i\}$  such that  $e_i = \frac{a_i}{b_i}$ , so there is  $a'_i \in \{0, 1, \dots, b\}$  such that  $e_i = \frac{a'_i}{b}$ , so  $e_i \in \mathcal{M}_b$ . Moreover,  $e_n = 1 - \sum_{i \in N \setminus \{n\}} e_i = \frac{b - \sum_{i \in N \setminus \{n\}} a'_i}{b}$ , so  $e_n \in \mathcal{M}_b$ . Altogether, then,  $c \leq b \leq b_1^{n-1} = p^{n-1}$ , as desired.

For the final item, assume by way of contradiction we have such a function  $f$ . For each  $m \in \{3, 4, \dots\}$ , let  $e^m \in E$  be such that one agent has entitlement  $\frac{1}{2} - \frac{1}{2m+1}$ , another agent has entitlement  $\frac{1}{2} + \frac{1}{2m+1}$ , and all other agents have entitlement zero. Then for each  $m \in \{3, 4, \dots\}$ , we have  $F(e^m) = 3$  and  $P(e^m) = 4m + 2$ . But then there is  $m \in \{3, 4, \dots\}$  such that  $P(e^m) > f(3) = f(F(e^m))$ , contradicting that  $f(F(e^m)) \geq P(e^m)$ . ■

Using [Proposition 2](#), we can translate [Proposition 1](#) from a lower bound involving precision to one involving clonage.

**THEOREM 1 (MAIN RESULT):** Fix a setting with at least two agents. For each clonage level  $c \in \mathbb{N}$ ,  $C(e) > c$  implies  $\text{cost}(e) > \lfloor \log_2 \log_2 2c^{\frac{1}{n-1}} \rfloor$ .

If there is one agent, of course, then for the unique member of  $E$  we have  $\text{cost}(e) = 0$ . We conclude this section with our titular corollary about the irrationally entitled.

**COROLLARY 1:** Fix a setting and an entitlement profile  $e$ . If there is an agent whose entitlement is irrational, then no  $e$ -proportional protocol is bounded.

## 4.2 The meta-problem

In this section, we use [Theorem 1](#) to analyze how cost varies as the problem varies. To do so, we consider the meta-problem of selecting a protocol for each problem, also known as the algorithm selection problem ([Rice, 1976](#)). The distinction between the mediator's problem and the meta-problem is worth emphasizing.

- A *problem* is given by a setting and an entitlement profile  $e$ . In this case, an *instance* is a kitchen measure profile  $\mu$ , whose associated *solutions* are the  $(e|\mu)$ -proportional allocations, and a *protocol* maps each instance to one of its solutions.
- A *meta-problem* is given by a setting. In this case, a *meta-instance* is an associated problem given by an entitlement profile  $e$ , whose associated *meta-solutions* are the  $e$ -proportional protocols, and (iii) a *meta-protocol* maps each meta-instance to one of its meta-solutions.

We note that our methodology for analyzing the meta-problem, which is standard in the cake division literature, involves parameterized complexity theory ([Downey and Fellows, 1999](#); [Flum and Grohe, 2006](#)). In particular, in the terminology of [Flum and Grohe \(2006\)](#), a parameterization of a problem is a function that associates each instance with a parameter. We take the clonage index as our parameterization of the meta-problem, while other papers on cake division take the number of agents as the parameterization.

For the mediator's problem, which protocols are optimal depends on the performance metric, and we focused on a single metric: the worst-case communication cost measured in queries. For the meta-problem, the optimal meta-protocols are fixed as those that associate each problem with an optimal protocol, and we measure how their performance varies with clonage using two metrics.

**DEFINITION:** Fix a setting. Let  $\text{infcost} : \mathbb{N} \rightarrow \mathbb{R}_+ \cup \{\infty\}$  and  $\text{supcost} : \mathbb{N} \rightarrow \mathbb{R}_+ \cup \{\infty\}$  be defined such that for each clonage level  $c \in \mathbb{N}$ ,

$$\begin{aligned}\text{infcost}(c) &\equiv \inf\{\text{cost}(e) \mid e \in E \text{ and } C(e) = c\}, \text{ and} \\ \text{supcost}(c) &\equiv \sup\{\text{cost}(e) \mid e \in E \text{ and } C(e) = c\}.\end{aligned}$$

Recall that for each  $f : \mathbb{N} \rightarrow \mathbb{R}_+ \cup \{\infty\}$  and each  $g : \mathbb{N} \rightarrow \mathbb{R}_+$ , we use  $g$  to express the asymptotic order of  $f$  by writing

- $f(x) = \mathcal{O}(g(x))$  if there are  $x^* \in \mathbb{N}$  and  $M^* \in \mathbb{R}_+ \in [0, \infty)$  such that for each  $x \in \{x^*, x^* + 1, \dots\}$ ,  $f(x) \leq M^* \cdot g(x)$ , and
- $f(x) = \Omega(g(x))$  if there are  $x^* \in \mathbb{N}$  and  $M^* \in \mathbb{R}_+ \in [0, \infty)$  such that for each  $x \in \{x^*, x^* + 1, \dots\}$ ,  $f(x) \geq M^* \cdot g(x)$ .

We are interested in expressing the asymptotic order of **infcost** and **supcost**.

[Theorem CF1](#) and [Theorem CF2](#) together bound **supcost** both above and below, and our results do not improve either bound.

**COROLLARY CF** ([CSEH AND FLEINER, 2020](#)): Fix a setting. We have  $\mathbf{supcost}(c) = \mathcal{O}(\log_2 c)$  and  $\mathbf{supcost}(c) = \Omega(\log_3 c) = \Omega(\log_2 c)$ .

Instead, we provide a lower bound on **infcost**.

**COROLLARY 2**: Fix a setting. We have  $\mathbf{infcost}(c) = \Omega(\log_2 \log_2 c)$ .

Together, [Corollary CF](#) and [Corollary 2](#) allow us to formalize our key insight that clonage causes communication costs: a sufficient increase in clonage guarantees that costs increase.

**COROLLARY 3**: Fix a setting. For each clonage level  $c \in \mathbb{N}$ , there is clonage level  $c' \in \mathbb{N}$  with  $c' > c$  such that for each pair  $e, e' \in E$ ,  $C(e) \leq c$  and  $C(e') \geq c'$  implies  $\mathbf{cost}(e) < \mathbf{cost}(e')$ .

### 4.3 Proof discussion

Our main result follows from [Proposition 1](#) and [Proposition 2](#), and we have already proven the latter. In this section, we discuss the proof of the former. We begin by introducing the concepts needed to state our key lemma, [Lemma 3](#).

To begin, sometimes two records are equally informative in the sense that they are associated with the same restricted domain. It is sometimes convenient to use a record with a small collection of servings, and it is sometimes convenient to use a record with a large number of servings. In our analysis, we restrict attention to the case that the small collection consists of the cells of a partition and the large collection is the associated subalgebra.

**DEFINITION**: Fix a setting. A (finite) record  $(\mathcal{S}', \alpha')$  is a *(finite) partition record* if  $\mathcal{S}'$  is a finite partition. In this case, we often emphasize that the record is a partition record by writing  $(\mathcal{P}, \alpha)$ . We let  $\mathcal{A}_{\mathcal{P}} \subseteq \mathcal{S}$  denote the subalgebra of  $\mathcal{S}$  generated by  $\mathcal{P}$ , and for each  $i \in N$ , we let  $\alpha_i^\uparrow$  denote the *subalgebra extension of  $\alpha_i$  (to  $\mathcal{A}_{\mathcal{P}}$ )*: the function  $\alpha_i^\uparrow : \mathcal{A}_{\mathcal{P}} \rightarrow [0, 1]$  such that for each  $S \in \mathcal{A}_{\mathcal{P}}$  we have  $\alpha_i^\uparrow(S) = \sum_{\{S' \in \mathcal{P} \mid S' \subseteq S\}} \alpha_i(S')$ .

We can always use a partition record as a concise upper bound on the information available to the mediator, in which case we conveniently have an equivalent subalgebra record.

**OBSERVATION**: Fix a setting and let  $(\mathcal{S}', \alpha')$  be a (finite) record. For each record  $(\mathcal{P}, \alpha)$  such that (i)  $\mathcal{P}$  is the collection of atoms in the subalgebra of  $\mathcal{S}$  generated by  $\mathcal{S}'$ , and (ii) for each  $i \in N$  and each  $S' \in \mathcal{S}'$ ,  $\alpha_i(S') = \alpha_i^\uparrow(S')$ , we have that  $(\mathcal{P}, \alpha)$  is a partition record that is at least as informative as  $(\mathcal{S}', \alpha')$ .

OBSERVATION:<sup>11</sup> Fix a setting. For each (finite) partition record  $(\mathcal{P}, \alpha)$ , the record  $(\mathcal{A}_{\mathcal{P}}, \alpha^\uparrow)$  is equally informative:  $\mathbb{D}(\mathcal{P}, \alpha) = \mathbb{D}(\mathcal{A}_{\mathcal{P}}, \alpha^\uparrow)$ .

Our key idea is indexing the difficulty that the mediator faces at an information set using the number of cake clones that he requires to solve the problem with no further queries. For our purposes, it suffices to formalize this difficulty index only for two-agent settings and partition records. To avoid confusion with agent clones, we refer to copies of the cake as *replicas*, and we refer to the difficulty that the mediator faces as the partition record's *deficiency level*.

DEFINITION: Fix a setting with agents in  $N = \{1, 2\}$  and an entitlement profile  $e \in E$ . For each partition record  $(\mathcal{P}, \alpha)$  and each level  $\ell \in \mathbb{N}$ , we define the following.

- For each  $r \in \mathbb{N}$ , an  $(r|\mathcal{P})$ -hyperserving is a member of  $\mathcal{S}_{r|\mathcal{P}} \equiv (\mathcal{A}_{\mathcal{P}})^r$ . The suggested interpretation is that the cake is replaced by  $r$  identical replicas, with the property that each serving in a replica is worth  $\frac{1}{r}$  of the associated serving in the original cake, and the hyperserving  $(S_1, S_2, \dots, S_r)$  consists of a serving from each replica that is a union of partition cells.
- An  $(\ell|\mathcal{P})$ -hyperallocation is a member of  $\mathcal{X}_{\ell|\mathcal{P}} \equiv \cup_{r \in \{1, 2, \dots, \ell\}} \mathcal{S}_{r|\mathcal{P}}$ . For each  $X \in \mathcal{X}_{\ell|\mathcal{P}}$ , we define the *number of replicas for  $X$* ,  $r(X)$ , to be the unique  $r \in \{1, 2, \dots, \ell\}$  such that  $X \in \mathcal{S}_{r|\mathcal{P}}$ . The interpretation is that the cake is replaced by  $r(X)$  replicas, then agent 1 consumes  $X$  while agent 2 consumes the rest.
- For each  $i \in N$  we define the *hyperallocation extension of  $\alpha_i^\uparrow$  (to  $\mathcal{X}_{\ell|\mathcal{P}}$ )*: the function  $\alpha_i^{\uparrow\uparrow} : \mathcal{X}_{\ell|\mathcal{P}} \rightarrow [0, 1]$  such that for each  $X \in \mathcal{X}_{\ell|\mathcal{P}}$  we have  $\alpha_i^{\uparrow\uparrow}(X) \equiv \sum_{j=1}^{r(X)} \frac{\alpha_i^\uparrow(X_j)}{r(X)}$ .
- We say that  $(\mathcal{P}, \alpha)$  is  $\ell$ -deficient if for each  $X \in \mathcal{X}_{\ell|\mathcal{P}}$ , either  $\alpha_1^{\uparrow\uparrow}(X) < e_1$  or  $\alpha_2^{\uparrow\uparrow}(X) > e_1$ .

Observe that if  $(\mathcal{P}, \alpha)$  is 1-deficient, then no  $e$ -proportional protocol selects an allocation at an information set whose record is at most as informative as  $(\mathcal{P}, \alpha)$ .

As we use partition records as upper bounds for the information that the mediator currently has, so too do we use *ultraresponses* as upper bounds for the information that the mediator acquires.

DEFINITION: Fix a setting. For each valid partition record  $(\mathcal{P}, \alpha)$ , each query  $\mathbf{q} = (i, \kappa, S, p)$ , and each  $\mu \in \mathbb{D}(\mathcal{P}, \alpha)$ , we define the  $(\mathbf{q}|\mu, \mathcal{P}, \alpha)$ -ultraresponse,  $r(\mathbf{q}|\mu, \mathcal{P}, \alpha)$ , to be the partition record  $(\mathcal{P}^+, \alpha^+)$  such that (i)  $\mathcal{P}^+$  is the collection of atoms in the subalgebra of  $\mathcal{S}$  generated by  $\mathcal{P} \cup \{S, S_{\mathbf{q}|\mu_i}\}$ , and (ii) for each  $j \in N$ ,  $\alpha_j^+ = \mu_j \upharpoonright_{\mathcal{P}^+}$ . We denote the collection of these ultraresponses by  $\mathcal{R}(\mathbf{q}|\mathcal{P}, \alpha) \equiv \{r(\mathbf{q}|\mu, \mathcal{P}, \alpha) | \mu \in \mathbb{D}(\mathcal{P}, \alpha)\}$ .

We want to formalize the idea that if the mediator finds himself in a particularly difficult situation, then for each query he might ask, there is a response that leaves him

<sup>11</sup>We caution that in general, if we are given a (finite) record  $(\mathcal{S}', \alpha')$  and seek an equally informative record with more entries, then by iteratively adding uninformative new entries for complements and disjoint unions, we reach a record for the *Dynkin system* generated by  $\mathcal{S}'$ , and a record for the *subalgebra* generated by  $\mathcal{S}'$  may be strictly more informative. Consider, for example, the interval cake and  $\mathcal{S}' = \{(\frac{1}{5}, \frac{3}{5}], (\frac{2}{5}, \frac{4}{5}]\}$ : an associated appraisal generally does not allow us to deduce the value of  $(\frac{2}{5}, \frac{3}{5}]$ .

in a situation that is still rather difficult. Our key lemma establishes this for the special case that the mediator’s query involves a cell of the current partition.

**LEMMA 3:** Fix a setting with two agents in  $N = \{1, 2\}$  and let  $e \in E$ . For each level  $\ell$ , each valid partition record  $(\mathcal{P}, \alpha)$  that is  $\ell$ -deficient, and each  $\mathbf{q} \in \mathcal{Q}$  such that the serving argument of  $\mathbf{q}$  belongs to  $\mathcal{P}$ , there is  $(\mathcal{P}^+, \alpha^+) \in \mathcal{R}(\mathbf{q}|\mathcal{P}, \alpha)$  that is  $\lfloor \sqrt{\ell/2} \rfloor$ -deficient.

**PROOF SKETCH:** The proof (see [Appendix 1](#)) is largely constructive. In response to a query about a cell  $B$ , one of the agents cuts  $B$  into two sub-cells  $B'$  and  $B''$ , resulting in  $\mathcal{P}^+$ , and these details do not matter. All that matters is one number: how much of the value of  $B$  does the second agent assign to  $B'$ ? The goal is to find a value  $v$  answering this question such that for  $\ell^+ \equiv \lfloor \sqrt{\ell/2} \rfloor$ , whenever the first agent considers an  $(\ell^+|\mathcal{P}^+)$ -hyperallocation to be a candidate solution, the second agent rules it out.

Let us say an  $(\ell^+|\mathcal{P}^+)$ -hyperallocation is  $B'$ -heavy if the first agent receives more copies of  $B'$  than  $B''$ ,  $B''$ -heavy if the first agent receives more copies of  $B''$  than  $B'$ , and balanced otherwise. The balanced  $(\ell^+|\mathcal{P}^+)$ -hyperallocations are ruled out because they are equivalent to  $(\ell^+|\mathcal{P})$ -hyperallocations and  $(\mathcal{P}, \alpha)$  is  $\ell$ -deficient, high values of  $v$  rule out  $B'$ -heavy  $(\ell^+|\mathcal{P}^+)$ -hyperallocations, and low values of  $v$  rule out  $B''$ -heavy  $(\ell^+|\mathcal{P}^+)$ -hyperallocations. The problem is therefore only interesting if there are both  $B'$ -heavy candidates and  $B''$ -heavy candidates, as in this case we cannot simply rule out the only candidates with an extreme response from the second agent.

The key argument is that in this case, the largest value that accepts a  $B'$ -heavy candidate is less than the smallest value that accepts a  $B''$ -heavy candidates, so that any intermediate value will do; this is Step 7 of the proof. Indeed, otherwise there are a value  $v$ , a  $B'$ -heavy candidate  $X'$ , and a  $B''$ -heavy candidate  $X''$  such that with  $v$  the second agent declares both  $X'$  and  $X''$  acceptable. In this case, however, we show that by appropriately replicating  $X'$  and  $X''$ , we can construct an  $(\ell|\mathcal{P}^+)$ -hyperallocation that solves the problem and is equivalent to an  $(\ell|\mathcal{P})$ -hyperallocation, contradicting that  $(\mathcal{P}, \alpha)$  is  $\ell$ -deficient.

From here, [Lemma 4](#) extends [Lemma 3](#) to arbitrary queries, for which the mediator’s selected serving and the cutter’s response together divide each existing cell into up to three sub-cells, as in [Cseh and Fleiner \(2020\)](#). Loosely, the idea is to assign negligible value to sub-cells so that it is almost as though one cell was divided into two sub-cells, so that we can derive the desired conclusion from [Lemma 3](#); see [Appendix 2](#) for the details.

We complete the proof of [Proposition 1](#) in [Appendix 3](#). First, we use [Lemma 4](#) to establish the proposition for the two-agent case. Intuitively, if the entitlement profile has a lot of precision, then the initial information set has a lot of deficiency, so since bad responses can always keep deficiency relatively intact, an unlucky mediator may face a long sequence of bad responses and remain unable to solve the problem. The one-agent case is trivial, so to conclude we handle the case of more than two agents using the two-agent result. Intuitively, we cannot have a protocol whose cost is too low relative to the precision for more than two agents, because then could imitate that protocol for one of the highest-precision agents in a two-agent setting where the other entitlements are merged, contradicting our two-agent result. We emphasize that this argument relies on precision instead of clonage because precision is defined at the agent level before it is aggregated for the group, but clonage is defined only at the group level; this is why we complete the proof of [Theorem 1](#) by separately applying [Proposition 2](#).

## 5 Further results

### 5.1 Preference foundation for kitchen measures

What, exactly, are we assuming about an agent's preferences when we assume they have a kitchen measure representation? We provide an axiomatic answer to this question by modifying a classic result of [Savage \(1954\)](#), replacing one of Savage's technical axioms with two new axioms involving the knife block.

To begin, an agent's preferences are used to compare servings, and we assume that they satisfy some basic conditions that are necessary to admit a probability measure representation. In particular, we assume that the preferences form a *qualitative probability* ([Bernstein, 1917](#); [de Finetti, 1937](#); [Koopman, 1940](#)), an object originally interpreted as beliefs about the relative likelihood of uncertain events.

DEFINITION: Fix a kitchen and let  $\succsim$  be a binary relation on  $\mathcal{S}$ . We say that  $\succsim$  satisfies

- *order* if  $\succsim$  is complete and transitive;
- *separability* if for each triple  $A, B, S \in \mathcal{S}$  such that  $A \cap S = B \cap S = \emptyset$ , we have  $A \succsim B$  if and only if  $A \cup S \succsim B \cup S$ ;
- *monotonicity* if for each pair  $A, B \in \mathcal{S}$ ,  $A \subseteq B$  implies  $B \succsim A$ ; and
- *non-degeneracy* if there are  $A, B \in \mathcal{S}$  such that  $A \succ B$ .

We say that  $\succsim$  is a *qualitative probability* if it satisfies all of these axioms.

Of these axioms, the most substantial assumption for cake division is separability. Indeed, this implies that given three equivalent slices, two chocolate and one vanilla, any union of a pair of these slices is equivalent: chocolate-vanilla is no better or worse than chocolate-chocolate.

There are qualitative probabilities without probability measure representations ([Kraft, Pratt, and Seidenberg, 1959](#)). In order to guarantee probability measure representation, [Savage \(1954\)](#) introduced the following technical axioms.

DEFINITION: Fix a kitchen. We say that a qualitative probability  $\succsim$  satisfies

- *fineness* if for each  $A \in \mathcal{S}$  such that  $A \succ \emptyset$ , there are  $m \in \mathbb{N}$  and  $(B_t)_{t \in \{1, 2, \dots, m\}} \in \mathcal{S}^m$  such that (i)  $\{B_1, B_2, \dots, B_m\}$  is a partition of  $C$ , and (ii) for each  $t \in \{1, 2, \dots, m\}$ , we have  $A \succsim B_t$ ; and
- *tightness* if for each pair  $A, B \in \mathcal{S}$ , if (i) for each  $A' \in \mathcal{S}$  such that  $A' \succ \emptyset$  and  $A \cap A' = \emptyset$ , we have  $A \cup A' \succsim B$ , and (ii) for each  $B' \in \mathcal{S}$  such that  $B' \succ \emptyset$  and  $B \cap B' = \emptyset$ , we have  $B \cup B' \succsim A$ , then we have that  $A \sim B$ .

Savage's result was originally stated for qualitative probabilities on power sets, and the proof applies for qualitative probabilities on  $\sigma$ -algebras, but our model involves qualitative probabilities on algebras. We are therefore interested in the following generalization of Savage's theorem.<sup>12</sup>

<sup>12</sup>For alternate axioms for qualitative probabilities on algebras, see [Chateauneuf \(1985\)](#) and [Barbanel and Taylor \(1995\)](#).

THEOREM WM ([WAKKER, 1981](#); [MARINACCI, 1993](#)): Fix a kitchen. If a qualitative probability  $\succsim$  satisfies *fineness* and *tightness*, then there is a probability measure  $\mu_0$  such that for each pair  $A, B \in \mathcal{S}$ , we have  $A \succsim B$  if and only if  $\mu_0(A) \geq \mu_0(B)$ . In this case,  $\mu_0$  is the unique such kitchen measure, and the range of  $\mu_0$  is a dense subset of the unit interval.

We seek the stronger conclusion that preferences are represented by a kitchen measure, and accordingly introduce two axioms that correspond to the two conditions for kitchen measures.

DEFINITION: Fix a kitchen. We say that a qualitative probability  $\succsim$  satisfies

- *sliver nullity* if for each  $S \in \mathcal{S}$ , we have  $S \sim \emptyset$  if and only if  $S \in \mathcal{S}^\circ$ ; and
- *knife continuity* if for each  $\kappa \in \mathbb{K}$  and each  $A \in \mathcal{S}$ , both  $\{t \in [0, 1] \mid A \succsim \kappa_t\}$  and  $\{t \in [0, 1] \mid \kappa_t \succsim A\}$  are closed.

The first axiom requires no explanation. The second axiom requires that each knife is a continuous function from the unit interval to the collection of servings when the latter is associated with the coarsest topology for which upper contour sets and lower contour sets—that is, sets of the form  $\{B \in \mathcal{S} \mid B \succsim A\}$  and  $\{B \in \mathcal{S} \mid A \succsim B\}$ , respectively—are closed. In other words, each knife is a path from the empty serving to the entire cake, and moreover one that is monotonic with respect to set inclusion.

It follows from [Theorem WM](#) that Savage’s axioms and our new axioms are together sufficient for kitchen measure representation, and it is not too difficult to show that they are also necessary. Less obvious, however, is that tightness is implied by the other axioms.

THEOREM 2: Fix a kitchen. A qualitative probability  $\succsim$  satisfies *fineness*, *sliver nullity*, and *knife continuity* if and only if there is a kitchen measure  $\mu_0$  such that for each pair  $A, B \in \mathcal{S}$ , we have  $A \succsim B$  if and only if  $\mu_0(A) \geq \mu_0(B)$ . In this case,  $\mu_0$  is the unique such kitchen measure.

The proof is in [Appendix 4](#).

## 5.2 The mediator and the adversary

We conclude by investigating some strategic considerations raised by our formalization of protocols as pure strategies.

We begin by discussing a game for which the strategic analysis is not so interesting. Recall that in the game from [Section 3.4](#), an indifferent nature makes an unobserved choice of measure profile and thereafter the mediator is the only player. If we replace nature with an *adversary* whose ranking of outcomes is the opposite of the mediator’s, then we can directly apply our analysis of the optimal protocol with a different interpretation: we analyze the strategy that is optimal for the mediator when the adversary selects his strategy after observing the mediator’s. Unfortunately, repeating this exercise for the adversary yields little guidance: all strategy profiles are equivalent, because after observing the adversary’s choice of measure profile, the mediator will immediately produce an  $e$ -proportional allocation. Altogether, we have an example of a two-player zero-sum

game with neither a (pure strategy) value (von Neumann, 1928) nor a (pure strategy) Nash equilibrium (Nash, 1951). This is a familiar phenomenon; a classic example is Matching Pennies.

There is, however, a natural variant of the preceding game for which the strategic analysis is more interesting. In particular, consider the perfect information game where iteratively the mediator poses queries and the adversary responds.

DEFINITION: Fix a setting. For each  $e \in E$ , the *adversary game given  $e$*  is the two player game defined as follows.

- There are two players: the *mediator* and the *adversary*. There are also histories partially ordered by precedence, as well as active players and actions associated with non-terminal histories and payoffs. The precedence order is such that each history and action at that history together determine an *immediate successor*, or an earliest history that follows. We define these remaining ingredients below, defining precedence entirely using immediate successors. The game has perfect information: all information sets are singletons.
- A history is *odd* if it is non-terminal with an even number of predecessors, and *even* if it is non-terminal with an odd number of predecessors. The initial history is odd. At each odd history, the active player is the mediator and the set of available actions is  $\mathcal{Q} \cup \mathcal{X}$ . If the mediator selects a query, then the immediate successor is non-terminal. If the mediator selects an allocation, then the immediate successor is terminal.
- At each even history, the active player is the adversary. Necessarily, the previous action was some query  $\mathbf{q}$ , and the adversary must select a response: the set of available actions is  $\{\mathbf{r}(\mathbf{q}|\mu) | \mu \in \mathbb{D}\}$ . Regardless of the adversary's action, the next history is non-terminal.
- Each odd history  $h$  is identified with the finite sequence of actions that leads to  $h$  from the initial history: there is a (*query*) *count*  $c \in \mathbb{N}_0$  such that  $h = (\mathbf{q}_1, \mathbf{r}_1, \mathbf{q}_2, \mathbf{r}_2, \dots, \mathbf{q}_c, \mathbf{r}_c)$ . In this case, the  *$h$ -domain* is  $\mathbb{D}(h) \equiv \cap_{t \in \{1, 2, \dots, c\}} \mathbb{D}(\mathbf{r}_t)$ , where by convention  $c = 0$  implies  $\mathbb{D}(h) = \mathbb{D}$ .
- A play is a maximal chain of histories. If a play ends at a terminal history after odd history  $h$  with query count  $c$ , and if the allocation  $X$  selected by the mediator has the property that for each  $\mu \in \mathbb{D}(h)$  we have that  $X$  is  $(e|\mu)$ -proportional, then the mediator's *payoff* is  $-c$ ; otherwise it is  $-\infty$ . Observe that at an odd history  $h$  with query count  $c$ , if  $\mathbb{D}(h) = \emptyset$  because the adversary's responses are inconsistent, then the mediator can receive payoff  $-c$  by selecting any allocation. The adversary's *payoff* is the negative of the mediator's payoff.

For each player, a (*pure*) *strategy* is a function that associates each history at which he is the active player with an action available at that history. We let  $\Sigma_m$  denote the set of strategies for the mediator and let  $\Sigma_a$  denote the set of strategies for the adversary. Finally, for each strategy profile  $(\sigma_m, \sigma_a) \in \Sigma_m \times \Sigma_a$ , we let  $\mathbb{P}(\sigma_m, \sigma_a)$  denote the associated payoff to the mediator, which the mediator seeks to maximize and the adversary seeks to minimize.<sup>13</sup>

<sup>13</sup>We use *payoffs* instead of *utilities* to emphasize that we are not considering randomization: the payoffs are not derived from preferences over lotteries and we are not considering mixed strategies.

We are interested in analyzing the adversary games using the classic notions of value (von Neumann, 1928) and equilibrium (Nash, 1951).

DEFINITION: Fix a setting, let  $e \in E$ , and consider the associated adversary game.

- If  $\sup_{\sigma_m \in \Sigma_m} \inf_{\sigma_a \in \Sigma_a} \mathbb{P}(\sigma_m, \sigma_a) = \inf_{\sigma_a \in \Sigma_a} \sup_{\sigma_m \in \Sigma_m} \mathbb{P}(\sigma_m, \sigma_a)$ , then this is *the game's value*; otherwise *the game has no value*.
- A strategy profile  $(\sigma_m, \sigma_a)$  is a (*pure strategy*) *Nash equilibrium* if (i) for each  $\sigma'_m \in \Sigma_m$ ,  $\mathbb{P}(\sigma_m, \sigma_a) \geq \mathbb{P}(\sigma'_m, \sigma_a)$ , and (ii) for each  $\sigma'_a \in \Sigma_a$ ,  $\mathbb{P}(\sigma_m, \sigma_a) \leq \mathbb{P}(\sigma_m, \sigma'_a)$ .

In each adversary game, each player has an infinite set of strategies without any further mathematical structure that the payoff function is required to respect. While it is possible to guarantee that such a game has a value if the payoff function is *concave-convexlike* (Fan, 1953; Sion, 1958), that is not the case here, and we are aware of no general theorem that guarantees the adversary games have values. Fortunately, we can do so using our earlier analysis.

THEOREM 3: Fix a setting and let  $e \in E$ . If all entitlements are rational, then the game has a value and this value is achieved in a Nash equilibrium. If there is an irrational entitlement, then the game's value is  $-\infty$  but the game has no Nash equilibrium.

The proof is in Appendix 5. The first conclusion follows directly from the existence of a bounded protocol when all entitlements are rational (Steinhaus, 1948), but the second conclusion does not follow directly from our result that there is no bounded protocol when some entitlement is irrational (Corollary 1): for each cost we must construct an adversary strategy that guarantees that cost, and for this we require our deeper insights about deficient partition records and ultraresponses (Lemma 4).

Theorem 3 reveals that so far as value and equilibrium are concerned, any adversary game with an irrational entitlement is effectively a one-player game: the mediator is effectively a dummy player and the adversary can effectively select any cost  $c \in \mathbb{N}$  to transfer to himself from the mediator. In this way, we can articulate the idea that this is a hard problem for the mediator using classic notions from game theory without appealing to randomization.

## Appendix 1

In this appendix, we prove Lemma 3. To do so, we first prove Lemma 2, which provides a convenient description of when a record is an ultraresponse.

LEMMA 2: Fix a setting. For each valid partition record  $(\mathcal{P}, \alpha)$ , each query  $\mathbf{q} = (i, \kappa, S, p)$ , and each record  $(\mathcal{P}^+, \alpha^+)$ ,  $(\mathcal{P}^+, \alpha^+) \in \mathcal{R}(\mathbf{q}|\mathcal{P}, \alpha)$  if and only if it satisfies the following conditions.

- *Cut compatibility.* There is  $\mu \in \mathbb{D}(\mathcal{P}, \alpha)$  such that (i)  $\mathcal{P}^+$  is the collection of atoms in the subalgebra of  $\mathcal{S}$  generated by  $\mathcal{P} \cup \{S, S_{\mathbf{q}|\mu_i}\}$ , and (ii)  $\alpha_i^{+\uparrow}(S_{\mathbf{q}|\mu_i}) = p \cdot \alpha_i^{+\uparrow}(S)$ .
- *Appraisal compatibility.* For each  $j \in N$  and each  $A \in \mathcal{P}$ ,  $\sum_{\{B \in \mathcal{P}^+ | B \subseteq A\}} \alpha_j^+(B) = \alpha_j(A)$ .

- *Sliver compatibility.* For each  $j \in N$  and each  $A \in \mathcal{P}^+$ ,  $\alpha_j^+(A) = 0$  if and only if  $A \in \mathcal{S}^\circ$ .

PROOF: It is straightforward to verify that the conditions are necessary using [Lemma 1](#); we omit the argument. To see that the conditions are sufficient, assume the hypotheses and let  $\mu$  be the member of  $\mathbb{D}(\mathcal{P}, \alpha)$  promised by cut compatibility.

First, we construct  $\mu^* \in \mathbb{D}$ . Indeed, fix  $j \in N$ . We define  $\mu_j^* : \mathcal{S} \rightarrow [0, 1]$  to be such that for each  $A \in \mathcal{S}$ ,  $\mu_j^*(A) \equiv \sum_{B \in \mathcal{P}^+ \setminus \mathcal{S}^\circ} \alpha_j^+(B) \cdot \frac{\mu_j(A \cap B)}{\mu_j(B)}$ , which for emphasis does not involve division by zero because  $B \in \mathcal{P}^+ \setminus \mathcal{S}^\circ$  implies  $\mu_j(B) > 0$ . By null slivers, appraisal compatibility, and the fact that  $(\mathcal{P}, \alpha)$  is valid, we have  $\sum_{B \in \mathcal{P}^+ \setminus \mathcal{S}^\circ} \alpha_j^+(B) = \sum_{B \in \mathcal{P}^+} \alpha_j^+(B) = \sum_{A \in \mathcal{P}} \alpha_j(A) = 1$ , so  $\mu_j^*$  is a convex combination of probability measures that satisfy knife divisibility, from which it is straightforward to show that  $\mu_j^*$  is a probability measure that satisfies knife divisibility. Moreover, since  $\mu_j$  satisfies null slivers and  $\alpha_j^+$  satisfies sliver compatibility, thus  $\mu_j^*$  satisfies null slivers. Altogether, then,  $\mu_j^*$  is a kitchen measure. Since  $j \in N$  was arbitrary, thus  $\mu^* \in \mathbb{D}$ , as desired.

To conclude, we prove that  $\mu^* \in \mathbb{D}(\mathcal{P}, \alpha)$ . To begin, by definition of  $\mu_i^*$ , the fact that each member of  $\mathcal{P}^+$  is either contained in  $S_{\mathbf{q}|\mu_i}$  or  $S \setminus S_{\mathbf{q}|\mu_i}$ , sliver compatibility, and the definition of subalgebra extension, respectively, we have

$$\begin{aligned} \mu_i^*(S_{\mathbf{q}|\mu_i}) &= \sum_{B \in \mathcal{P}^+ \setminus \mathcal{S}^\circ} \alpha_i^+(B) \cdot \frac{\mu_i(S_{\mathbf{q}|\mu_i} \cap B)}{\mu_i(B)} \\ &= \sum_{\{B \in \mathcal{P}^+ \setminus \mathcal{S}^\circ \mid B \subseteq S_{\mathbf{q}|\mu_i}\}} \alpha_i^+(B) \\ &= \sum_{\{B \in \mathcal{P}^+ \mid B \subseteq S_{\mathbf{q}|\mu_i}\}} \alpha_i^+(B) \\ &= \alpha_i^{+\uparrow}(S_{\mathbf{q}|\mu_i}). \end{aligned}$$

By the same argument we have  $\mu_i^*(S \setminus S_{\mathbf{q}|\mu_i}) = \alpha_i^{+\uparrow}(S \setminus S_{\mathbf{q}|\mu_i})$ , so by the additivity of both  $\mu_i^*$  and  $\alpha_i^{+\uparrow}$  we have  $\mu_i^*(S) = \alpha_i^{+\uparrow}(S)$ . Thus by cut compatibility, we have

$$\begin{aligned} \mu_i^*(S_{\mathbf{q}|\mu_i}) &= \alpha_i^{+\uparrow}(S_{\mathbf{q}|\mu_i}) \\ &= p \cdot \alpha_i^{+\uparrow}(S) \\ &= p \cdot \mu_i^*(S). \end{aligned}$$

Moreover, fix  $t \in [0, \tau(\mathbf{q}|\mu_i))$ . First, by definition of  $\tau(\mathbf{q}|\mu_i)$  we have  $S \cap \kappa_t \subseteq S_{\mathbf{q}|\mu_i}$ , so for each  $B \in \mathcal{P}^+$  we have  $\mu_i^*((S \cap \kappa_t) \cap B) \leq \mu_i^*(S_{\mathbf{q}|\mu_i} \cap B)$ . Second, by definition of  $\tau(\mathbf{q}|\mu_i)$  and [Lemma 1](#) we have  $\mu_i(S \cap \kappa_t) < p \cdot \mu_i(S) = \mu_i(S_{\mathbf{q}|\mu_i})$ , so there is  $B \in \mathcal{P}^+$  such that  $\mu_i((S \cap \kappa_t) \cap B) < \mu_i(S_{\mathbf{q}|\mu_i} \cap B)$  and thus  $\mu_i^*((S \cap \kappa_t) \cap B) < \mu_i^*(S_{\mathbf{q}|\mu_i} \cap B)$ . Thus by additivity of  $\mu_i^*$ , we have  $\mu_i^*(S \cap \kappa_t) < \mu_i^*(S_{\mathbf{q}|\mu_i}) = p \cdot \mu_i^*(S)$ . Since  $t \in [0, \tau(\mathbf{q}|\mu_i))$  was arbitrary, altogether we have  $S_{\mathbf{q}|\mu_i^*} = S_{\mathbf{q}|\mu_i}$ , from which it is straightforward to verify that  $(\mathcal{P}^+, \alpha^+) = \mathbf{r}(\mathbf{q}|\mu^*, \mathcal{P}, \alpha)$ . ■

We conclude this appendix by providing [Lemma 3](#).

LEMMA 3 (RESTATED): Fix a setting with two agents in  $N = \{1, 2\}$  and let  $e \in E$ . For each level  $\ell$ , each valid partition record  $(\mathcal{P}, \alpha)$  that is  $\ell$ -deficient, and each  $\mathbf{q} \in \mathcal{Q}$

such that the serving argument of  $\mathbf{q}$  belongs to  $\mathcal{P}$ , there is  $(\mathcal{P}^+, \alpha^+) \in \mathcal{R}(\mathbf{q}|\mathcal{P}, \alpha)$  that is  $\lfloor \sqrt{\ell/2} \rfloor$ -deficient.

PROOF: Let  $\ell$ ,  $(\mathcal{P}, \alpha)$ , and  $\mathbf{q}$  satisfy the hypotheses. Several concepts in this proof come in pairs: we use 1 and 2 to distinguish the agents, ' and '' to distinguish the two servings in  $\mathcal{P}^+ \setminus \mathcal{P}$ , and  $+$  to distinguish objects for the new partition from associated objects for the starting partition. To complete the proof, we select an arbitrary kitchen measure profile  $\mu$  from a nonempty set, then construct the desired  $(\mathcal{P}^+, \alpha^+)$ .

STEP 1: *Introduce concepts and notation to be used for the remaining steps.*

STEP 1A: *Introduce preliminary notation, then construct (i) the servings  $B'$  and  $B''$ , (ii) the partition  $\mathcal{P}^+$ , and (iii) the appraisal  $\alpha_1^+$ . Moreover, proceed under the assumption that  $B'$  and  $B''$  are not slivers.*

First, we introduce preliminary notation. Define  $\ell^+ \equiv \lfloor \sqrt{\ell/2} \rfloor$ , let  $s \equiv |\mathcal{P}|$  denote the number of slices in  $\mathcal{P}$ , and index the members of  $\mathcal{P}$  by  $\{A_j\}_{j \in \{1, 2, \dots, s-1\}} \cup \{B\}$  such that  $B$  is the serving argument of  $\mathbf{q}$ . Furthermore, denote all the arguments of  $\mathbf{q}$  by  $(1, \kappa, B, p_1)$ , so that (re-indexing if necessary)  $\mathbf{q}$  asks agent 1 to cut, and let 2 denote the other agent. Observe that such a re-indexing is without loss of generality even though our notion of hyperallocation does not treat agents symmetrically: for each level  $\ell'$ , a partition record is  $\ell'$ -deficient before re-indexing if and only if it is  $\ell'$ -deficient after re-indexing. Define  $v_1 \equiv p_1 \cdot \alpha_1(B)$ , so that  $\mathbf{q}$  effectively asks 1 to use  $\kappa$  to construct a subset of  $B$  worth  $v_1$ .

Second, we construct  $B'$ ,  $B''$ ,  $\mathcal{P}^+$ , and  $\alpha_1^+$ . To begin, since  $(\mathcal{P}, \alpha)$  is valid, thus we can select an arbitrary  $\mu \in \mathbb{D}(\mathcal{P}, \alpha)$ ; we remark that for the ultrareponse  $(\mathcal{P}^+, \alpha^+)$  that we will construct, both  $\mathcal{P}^+$  and  $\alpha_1^+$  will be consistent with  $\mu_1$ . We define  $B' \equiv S_{\mathbf{q}|\mu_1}$  and define  $B'' \equiv B \setminus B'$ . If  $B' \in \mathcal{S}^\circ$  or  $B'' \in \mathcal{S}^\circ$ , then for  $(\mathcal{P}^+, \alpha^+) \equiv \mathbf{r}(\mathbf{q}|\mu, \mathcal{P}, \alpha)$  we have  $\mathbb{D}(\mathcal{P}, \alpha) = \mathbb{D}(\mathcal{P}^+, \alpha^+)$ , from which the desired conclusion directly follows; thus let us assume  $B' \notin \mathcal{S}^\circ$  and  $B'' \notin \mathcal{S}^\circ$ . Define  $\mathcal{P}^+ \equiv (\mathcal{P} \setminus \{B\}) \cup \{B', B''\}$ , and observe that this is indeed a partition. Finally, define  $\alpha_1^+ \equiv \mu_1 \upharpoonright_{\mathcal{P}^+}$ . Observe that by construction, (i) for each  $A \in \mathcal{P}^+ \setminus \{B', B''\}$  we have  $\alpha_1^+(A) = \alpha_1(A)$ , (ii)  $\alpha_1^+(B') = v_1$ , and (iii)  $\alpha_1^+(B'') = \alpha_1^+(B) - v_1$ .

Finally, we highlight an abuse of notation: throughout this proof, we always use the same notation for an appraisal, its subalgebra extension, and its hyperallocation extension. For example, instead of distinguishing between  $\alpha_1^+$ ,  $\alpha_1^{+\uparrow}$ , and  $\alpha_1^{+\uparrow\uparrow}$ , we simply write  $\alpha_1^+$  for all three functions, with the particular function we are using determined not by notation but by context.

STEP 1B: *Introduce the possible values that 2 assigns to  $B'$ , their associated appraisals and records, and related concepts.*

First, we define the set of possible values that 2 assigns to  $B'$  (given that neither  $B'$  nor  $B''$  is a sliver), its supremum, and its closure. In particular, (i) define  $v^{\sup} \equiv \alpha_2(B)$ , (ii) define  $V \equiv (0, v^{\sup})$  to be the set of possible values that 2 assigns to  $B'$ , and (iii) define  $\bar{V} \equiv [0, v^{\sup}]$ . Since  $B'$  and  $B''$  are not slivers, thus  $B = B' \cup B''$  is not a sliver, so as  $(\mathcal{P}, \alpha)$  is valid we have that  $v^{\sup} = \alpha_2(B) > 0$ .

Second, for each  $v \in \bar{V}$  we define (i) the  $\mathcal{P}^+$ -appraisal  $\alpha_2^v : \mathcal{P}^+ \rightarrow [0, 1]$ , and (ii) the record  $\mathbf{r}_v$ , as follows:

- $\alpha_2^v$  is defined by (i) for each  $A \in \mathcal{P}^+ \setminus \{B', B''\}$ ,  $\alpha_2^v(A) = \alpha_2(A)$ , (ii)  $\alpha_2^v(B') = v$ , and (iii)  $\alpha_2^v(B'') = \alpha_2^v(B) - v$ , and
- $r_v \equiv (\mathcal{P}^+, (\alpha_1^+, \alpha_2^v))$ .

Observe that by Lemma 2, for each  $v \in V$  we have  $r_v \in \mathcal{R}(\mathbf{q}|\mathcal{P}, \alpha)$ . By contrast, for each  $v \in \{0, v^{\sup}\}$ ,  $r_v$  is not valid because it assigns zero to a serving that is not a sliver, and thus  $r_v \notin \mathcal{R}(\mathbf{q}|\mathcal{P}, \alpha)$ . Even so, it is convenient to use the values in  $\{0, v^{\sup}\}$  for some of our arguments.

STEP 1C: *Introduce weights, lines, and candidates.*

First, for each hyperallocation  $X \in \mathcal{X}_{\ell^+|\mathcal{P}^+}$ , we write the value that 2 assigns to  $X$  as a function of the value that he assigns to  $B'$ . In particular, for each  $X \in \mathcal{X}_{\ell^+|\mathcal{P}^+}$ , we define the weights  $w'(X), w''(X) \in \{0, 1, \dots, r(X)\}$  and the line  $\mathcal{L}_X : \bar{V} \rightarrow [0, 1]$ , as follows:

$$\begin{aligned} w'(X) &\equiv |\{j \in \{1, 2, \dots, r(X)\} | B' \subseteq X_j\}|, \\ w''(X) &\equiv |\{j \in \{1, 2, \dots, r(X)\} | B'' \subseteq X_j\}|, \text{ and} \\ \mathcal{L}_X(v) &\equiv \alpha_2^v(X) \\ &= \left[ \frac{[\sum_{j=1}^{r(X)} \alpha_2(X_j \setminus B)] + w''(X) \cdot v^{\sup}}{r(X)} + \left[ \frac{w'(X) - w''(X)}{r(X)} \right] \cdot v \right]. \end{aligned}$$

We say that  $X$  is (i) *balanced* if  $w'(X) = w''(X)$ , (ii)  *$B'$ -heavy* if  $w'(X) > w''(X)$ , and (iii)  *$B''$ -heavy* if  $w'(X) < w''(X)$ . Observe that a balanced  $X$  has a flat  $X$ -line, a  $B'$ -heavy  $X$  has an increasing  $X$ -line, and a  $B''$ -heavy  $X$  has a decreasing  $X$ -line.

Second, for each  $X \in \mathcal{X}_{\ell^+|\mathcal{P}^+}$ , we say that  $X$  is a *candidate* if  $\alpha_1^+(X) \geq e_1$ . To complete the proof, we construct  $v^* \in V$  such that for each candidate  $X$  we have  $\mathcal{L}_X(v^*) = \alpha_2^{v^*}(X) > e_1$ , as this directly implies that  $r_{v^*}$  is  $\ell^+$ -deficient. Intuitively, if 2 assigns such a value  $v^*$  to  $B'$ , then for each hyperallocation that is a candidate for being satisfactory in the sense that 1 measures his own hyperserving to be worth at least his own entitlement, the hyperallocation fails to be satisfactory because 2 measures 1's hyperserving to be worth more than 1's entitlement—and thus measures his own hyperserving to be worth less than his own entitlement.

STEP 2: *For each balanced candidate  $X^+ \in \mathcal{X}_{\ell^+|\mathcal{P}^+}$  and each  $v \in \bar{V}$ ,  $\mathcal{L}_{X^+}(v) > e_1$ .*

Let  $X^+$  and  $v$  satisfy the hypotheses. For convenience, let us write  $r^+ \equiv r(X^+)$ ,  $\alpha_2^+ \equiv \alpha_2^v$ ,  $w' \equiv w'(X^+)$ , and  $w'' \equiv w''(X^+)$ . Since  $X^+$  is balanced, thus  $w' = w''$ . Observe that  $r^+ \leq \ell^+ = \lfloor \sqrt{\ell/2} \rfloor \leq \ell$ .

Let  $X \in \mathcal{X}_{\ell|\mathcal{P}}$  with  $r(X) = r^+$  be such that for each  $j \in \{1, 2, \dots, r^+\}$ , (i)  $j \leq w'$  implies  $X_j \equiv (X_j^+ \setminus B) \cup B$ , and (ii)  $j > w'$  implies  $X_j \equiv X_j^+ \setminus B$ . Then for each  $i \in N$ ,

we have

$$\begin{aligned}
\alpha_i(X) &= \left[ \sum_{j=1}^{r^+} \frac{\alpha_i(X_j^+ \setminus B)}{r^+} \right] + w' \alpha_i(B) \\
&= \left[ \sum_{j=1}^{r^+} \frac{\alpha_i^+(X_j^+ \setminus B)}{r^+} \right] + w' \alpha_i^+(B') + w' \alpha_i^+(B'') \\
&= \left[ \sum_{j=1}^{r^+} \frac{\alpha_i^+(X_j^+ \setminus B)}{r^+} \right] + w' \alpha_i^+(B') + w'' \alpha_i^+(B'') \\
&= \alpha_i^+(X^+).
\end{aligned}$$

Since  $X^+$  is a candidate, thus we have  $\alpha_1(X) = \alpha_1^+(X^+) \geq e_1$ , so since  $(\mathcal{P}, \alpha)$  is  $\ell$ -deficient we have  $\mathcal{L}_{X^+}(v) = \alpha_2^+(X^+) = \alpha_2(X) > e_1$ , as desired.

STEP 3: For each  $B'$ -heavy candidate  $X^+ \in \mathcal{X}_{\ell^+|\mathcal{P}^+}$ ,  $\mathcal{L}_{X^+}(v^{\sup}) > e_1$ .

Let  $X^+$  satisfy the hypotheses. For convenience, let us write  $r^+ \equiv r(X^+)$ ,  $\alpha_2^+ \equiv \alpha_2^{v^{\sup}}$ ,  $w' \equiv w'(X^+)$ , and  $w'' \equiv w''(X^+)$ . Since  $X^+$  is  $B'$ -heavy, thus  $w' > w''$ . Observe that  $r^+ \leq \ell^+ = \lfloor \sqrt{\ell/2} \rfloor \leq \ell$ .

Let  $X \in \mathcal{X}_{\ell|\mathcal{P}}$  with  $r(X) = r^+$  be such that for each  $j \in \{1, 2, \dots, r^+\}$ , (i)  $j \leq w'$  implies  $X_j \equiv (X_j^+ \setminus B) \cup B$ , and (ii)  $j > w'$  implies  $X_j \equiv X_j^+ \setminus B$ . Then for each  $i \in N$ , we have

$$\begin{aligned}
\alpha_i(X) &= \left[ \sum_{j=1}^{r^+} \frac{\alpha_i(X_j^+ \setminus B)}{r^+} \right] + w' \alpha_i(B) \\
&= \left[ \sum_{j=1}^{r^+} \frac{\alpha_i^+(X_j^+ \setminus B)}{r^+} \right] + w' \alpha_i^+(B') + w' \alpha_i^+(B'') \\
&\geq \left[ \sum_{j=1}^{r^+} \frac{\alpha_i^+(X_j^+ \setminus B)}{r^+} \right] + w' \alpha_i^+(B') + w'' \alpha_i^+(B'') \\
&= \alpha_i^+(X^+).
\end{aligned}$$

Moreover, since  $\alpha_2^+(B'') = 0$ , thus the inequality above holds with equality for agent  $i = 2$ . Since  $X^+$  is a candidate, thus we have  $\alpha_1(X) \geq \alpha_1^+(X^+) \geq e_1$ , so since  $(\mathcal{P}, \alpha)$  is  $\ell$ -deficient we have  $\mathcal{L}_{X^+}(v^{\sup}) = \alpha_2^+(X^+) = \alpha_2(X) > e_1$ , as desired.

STEP 4: For each  $B''$ -heavy candidate  $X^+ \in \mathcal{X}_{\ell^+|\mathcal{P}^+}$ ,  $\mathcal{L}_{X^+}(0) > e_1$ .

The argument is the same as that for Step 3, except we use 0 instead of  $v^{\sup}$  and we transpose the roles of  $w'$  and  $w''$ .

STEP 5: Construct  $v', v'' \in \overline{V}$  such that

$$\begin{aligned}
v' &= \min\{v \in \overline{V} \mid \text{for each } B'\text{-heavy candidate } X, \mathcal{L}_X(v) \geq e_1\}, \text{ and} \\
v'' &= \max\{v \in \overline{V} \mid \text{for each } B''\text{-heavy candidate } X, \mathcal{L}_X(v) \geq e_1\},
\end{aligned}$$

and such that moreover  $v' < v^{\sup}$  and  $v'' > 0$ .

Before we begin, we observe that if there are a  $B'$ -heavy candidate and a  $B''$ -heavy candidate, then  $v'$  is the largest value that accepts a  $B'$ -heavy candidate and  $v''$  is the smallest value that accepts a  $B''$ -heavy candidate; thus these are the values mentioned in the final paragraph of the proof sketch.

First, we construct  $v'$ . If there is no  $B'$ -heavy candidate, then define  $v' \equiv 0$  and we are done; thus let us assume there is a  $B'$ -heavy candidate. In this case, by Step 3, for each  $B'$ -heavy candidate  $X$ ,  $\mathcal{L}_X$  is a strictly increasing line with  $\mathcal{L}_X(v^{\sup}) > e_1$ , so  $v(X) \equiv \min\{v \in \bar{V} \mid \mathcal{L}_X(v) \geq e_1\}$  is well-defined and less than  $v^{\sup}$ ; thus as there are finitely many  $B'$ -heavy candidates,  $v' \equiv \max\{v(X) \mid X \text{ is a } B'\text{-heavy candidate}\}$  is well-defined and satisfies the desired properties.

Second, we construct  $v''$ . The construction is analogous to the construction of  $v'$ , using  $B''$ -heavy candidates and Step 4 instead of  $B'$ -heavy candidates and Step 3; we omit the details.

STEP 6: *Proceed under the assumption that there are a  $B'$ -heavy candidate and a  $B''$ -heavy candidate.*

If there is no  $B'$ -heavy candidate and there is no  $B''$ -heavy candidate, then define  $v^* = \frac{v^{\sup}}{2}$ ; by Step 2, for each candidate  $X$  we have  $\mathcal{L}_X(v^*) > e_1$ , so we are done. If there is a  $B'$ -heavy candidate but there is no  $B''$ -heavy candidate, then define  $v^* = \frac{v' + v^{\sup}}{2}$ ; by Step 2 and Step 5, for each candidate  $X$  we have  $\mathcal{L}_X(v^*) > e_1$ , so we are done. If there is a  $B''$ -heavy candidate but there is no  $B'$ -heavy candidate, then define  $v^* = \frac{v''}{2}$ ; by Step 2 and Step 5, for each candidate  $X$  we have  $\mathcal{L}_X(v^*) > e_1$ , so we are done. Thus let us assume there are a  $B'$ -heavy candidate and a  $B''$ -heavy candidate.

STEP 7: *We have  $v' < v''$ .*

Assume for contradiction that  $v' \geq v''$  and let  $v^* \in (v'', v')$ . Using the definitions of  $v'$  and  $v''$ , it is easy to verify that there are  $B'$ -heavy candidate  $X' \in \mathcal{X}_{\ell+|\mathcal{P}^+}$  and  $B''$ -heavy candidate  $X'' \in \mathcal{X}_{\ell+|\mathcal{P}^+}$  such that  $\mathcal{L}_{X'}(v^*) \leq e_1$  and  $\mathcal{L}_{X''}(v^*) \leq e_1$ . For convenience, let us write  $r' \equiv r(X')$ ,  $r'' \equiv r(X'')$ , and  $\alpha_2^+ \equiv \alpha_2^{v^*}$ .

Intuitively, we construct a hyperallocation that (i) consists of some copies of  $X'$  and some copies of  $X''$ , (ii) involves a grand total of at most  $\ell$  replicas, and (iii) includes the same number of  $B'$  slices and  $B''$  slices. This  $(\ell|\mathcal{P}^+)$ -hyperallocation is equivalent to an  $(\ell|\mathcal{P})$ -hyperallocation that we refer to as  $X$ .

We begin by using  $X'$  and  $X''$  to define the weight lists  $(w'_S)_{S \in \mathcal{P}^+}$  and  $(w''_S)_{S \in \mathcal{P}^+}$  and the weight differences  $\delta'$  and  $\delta''$ , then using these notions to define the number of replicas  $r$  and the weight list  $(w_S)_{S \in \mathcal{P}^+}$ . In particular:

- for each  $S \in \mathcal{P}^+$ ,  $w'_S \equiv |\{j \in \{1, 2, \dots, r'\} \mid S \subseteq X'_j\}|$ ,
- for each  $S \in \mathcal{P}^+$ ,  $w''_S \equiv |\{j \in \{1, 2, \dots, r''\} \mid S \subseteq X''_j\}|$ ,
- $\delta' \equiv [w'_{B'} - w'_{B''}]$ ,
- $\delta'' \equiv [w''_{B''} - w''_{B'}]$ ,
- $r \equiv \delta'' r' + \delta' r''$ , and

- for each  $S \in \mathcal{P}^+$ ,  $w_S \equiv \delta'' w'_S + \delta' w''_S$ .

We use these objects to establish our contradiction. Intuitively, we construct a new hyperallocation that consists of  $\delta''$  copies of  $X'$  and  $\delta'$  copies of  $X''$ .<sup>14</sup>

Next, we make five observations. First, since  $r', r'' \in \{0, 1, \dots, \ell^+\}$ , thus for each  $S \in \mathcal{P}^+$  we have  $w'_S, w''_S \in \{0, 1, \dots, \ell^+\}$ . Second, we have  $\sum_{S \in \mathcal{P}^+} w'_S \alpha_2^+(S) = \mathcal{L}_{X'}(v^*) \leq e_1$  and  $\sum_{S \in \mathcal{P}^+} w''_S \alpha_2^+(S) = \mathcal{L}_{X''}(v^*) \leq e_1$ . Third, we have  $\delta', \delta'', r', r'' \in \{0, 1, \dots, \ell^+\}$  and thus  $r \in \{0, 1, \dots, 2(\ell^+)^2\}$ , and moreover  $\ell^+ = \lfloor \sqrt{\ell/2} \rfloor$ , so  $r \in \{0, 1, \dots, \ell\}$ . Fourth, for each  $S \in \mathcal{P}^+$ , since  $w'_S \in \{0, 1, \dots, r'\}$  and  $w''_S \in \{0, 1, \dots, r''\}$ , thus  $w_S \in \{0, 1, \dots, r\}$ . Fifth, we have

$$\begin{aligned} w_{B'} &= [w''_{B''} - w''_{B'}]w'_{B'} + [w'_{B'} - w'_{B''}]w''_{B'} \\ &= w''_{B''}w'_{B'} - w''_{B'}w'_{B''} \\ &= [w''_{B''} - w''_{B'}]w'_{B''} + [w'_{B'} - w'_{B''}]w''_{B''} \\ &= w_{B''}. \end{aligned}$$

Define  $w_B \equiv w_{B'} = w_{B''}$ .

Finally, for each  $j \in \{1, 2, \dots, r\}$ , define  $X_j \equiv \cup\{S \in \mathcal{P} | j \leq w_S\}$ . Since  $r \leq \ell$ , thus  $X \in \mathcal{X}_{\ell|\mathcal{P}}$ . Using the above definitions and observations, for each  $i \in N$  we have

$$\begin{aligned} \alpha_i(X) &= \sum_{j=1}^r \frac{\alpha_i(X_j)}{r} \\ &= \frac{\sum_{A \in \mathcal{P} \setminus \{B\}} w_A \alpha_i(A)}{r} + \frac{w_B \alpha_i(B)}{r} \\ &= \frac{\sum_{A \in \mathcal{P}^+ \setminus \{B', B''\}} w_A \alpha_i^+(A)}{r} + \frac{w_B \alpha_i^+(B)}{r} \\ &= \frac{\sum_{A \in \mathcal{P}^+ \setminus \{B', B''\}} w_A \alpha_i^+(A)}{r} + \frac{w_B \alpha_i^+(B') + w_B \alpha_i^+(B'')}{r} \\ &= \frac{\sum_{A \in \mathcal{P}^+ \setminus \{B', B''\}} w_A \alpha_i^+(A)}{r} + \frac{w_{B'} \alpha_i^+(B') + w_{B''} \alpha_i^+(B'')}{r} \\ &= \frac{\sum_{A \in \mathcal{P}^+} w_A \alpha_i^+(A)}{r} \\ &= \frac{\sum_{A \in \mathcal{P}^+} (\delta'' w'_A + \delta' w''_A) \alpha_i^+(A)}{r} \\ &= \delta'' \frac{\sum_{A \in \mathcal{P}^+} w'_A \alpha_i^+(A)}{r} + \delta' \frac{\sum_{A \in \mathcal{P}^+} w''_A \alpha_i^+(A)}{r} \\ &= \left(\frac{\delta'' r'}{r}\right) \frac{\sum_{A \in \mathcal{P}^+} w'_A \alpha_i^+(A)}{r'} + \left(\frac{\delta' r''}{r}\right) \frac{\sum_{A \in \mathcal{P}^+} w''_A \alpha_i^+(A)}{r''} \\ &= \left(\frac{\delta'' r'}{\delta'' r' + \delta' r''}\right) \alpha_i^+(X') + \left(\frac{\delta' r''}{\delta'' r' + \delta' r''}\right) \alpha_i^+(X''). \end{aligned}$$

But then (i) since  $X'$  and  $X''$  are candidates we have  $\alpha_1(X) \geq e_1$ , and (ii) since  $\alpha_2^+(X') = \mathcal{L}_{X'}(v^*) \leq e_1$  and  $\alpha_2^+(X'') = \mathcal{L}_{X''}(v^*) \leq e_1$  we have  $\alpha_2(X) \leq e_1$ , contradicting that  $(\mathcal{P}, \alpha)$  is  $\ell$ -deficient.

<sup>14</sup>We remark that at this point, an alternative path leads to a slightly tighter but less simple bound. In particular, if we let  $g$  denote the greatest common divisor of  $\delta'$  and  $\delta''$ , then we could complete an analogous argument by making only  $\frac{\delta''}{g}$  copies of  $X'$  and  $\frac{\delta'}{g}$  copies of  $X''$ . We choose to pursue the slightly looser but simpler bound.

STEP 8: *Conclude.*

By Step 7, we have  $v' < v''$ . Define  $v^* \equiv \frac{v' + v''}{2}$ . First, by Step 2, for each balanced candidate  $X$  we have  $\mathcal{L}_X(v^*) > e_1$ . Second, using the definition of  $v'$ , it is easy to verify that for each  $B'$ -heavy candidate  $X$  we have  $\mathcal{L}_X(v^*) > e_1$ . Finally, using the definition of  $v''$ , it is easy to verify that for each  $B''$ -heavy candidate  $X$  we have  $\mathcal{L}_X(v^*) > e_1$ . Altogether, then, for each candidate  $X$  we have  $\mathcal{L}_X(v^*) > e_1$ , so  $\mathbf{r}_{v^*}$  is  $\ell^+$ -deficient, as desired. ■

## Appendix 2

In this appendix, we prove [Lemma 4](#).

LEMMA 4 (RESTATEd): Fix a setting with two agents in  $N = \{1, 2\}$  and let  $e \in E$ . For each level  $\ell$ , each valid partition record  $(\mathcal{P}, \alpha)$  that is  $\ell$ -deficient, and each  $\mathbf{q} \in \mathcal{Q}$ , there is  $(\mathcal{P}^+, \alpha^+) \in \mathcal{R}(\mathbf{q}|\mathcal{P}, \alpha)$  that is  $\lfloor \sqrt{\ell/2} \rfloor$ -deficient.

PROOF: Let  $\ell$ ,  $(\mathcal{P}, \alpha)$ , and  $\mathbf{q}$  satisfy the hypotheses. To complete the proof, we select an arbitrary kitchen measure profile  $\mu$  from a nonempty set, then construct the desired  $(\mathcal{P}^+, \alpha^+)$ .

STEP 1: *Introduce preliminary notation and select the kitchen measure profile  $\mu$ , then construct the serving  $B'$  and the table  $\mathcal{T}$ .*

First, we introduce preliminary notation similar to that used in the proof of [Lemma 3](#). Define  $\ell^+ \equiv \lfloor \sqrt{\ell/2} \rfloor$ , let  $s \equiv |\mathcal{P}|$  denote the number of slices in  $\mathcal{P}$ , and index the members of  $\mathcal{P}$  by  $\{A_j\}_{j \in \{1, 2, \dots, s\}}$ . Denote the arguments of  $\mathbf{q}$  by  $(1, \kappa, B, p_1)$ , so that (re-indexing if necessary)  $\mathbf{q}$  asks agent 1 to cut, and let 2 denote the other agent; such a re-indexing is without loss of generality as argued at the start of the proof of [Lemma 3](#).

Second, we construct a table of servings, each given by a row and a column, and from this table we construct  $\mathcal{P}^+$ . We take  $\mathcal{P}$  to be the set of *rows*. Since  $(\mathcal{P}, \alpha)$  is valid, thus we can select an arbitrary  $\mu \in \mathbb{D}(\mathcal{P}, \alpha)$ ; we remark that for the ultraresponse  $(\mathcal{P}^+, \alpha^+)$  that we will construct,  $\mathcal{P}^+$  will be consistent with  $\mu_1$  while  $\alpha_1^+$  will be constructed by modifying  $\mu_1$ . We define  $B' \equiv S_{\mathbf{q}|\mu_1}$ , we define the set of *columns*  $\mathcal{C} \equiv \{B', B \setminus B', C \setminus B\}$ , and we define the *table*  $\mathcal{T} \equiv \{A \cap A' | A \in \mathcal{P}, A' \in \mathcal{C}\}$ .

Finally, we highlight an abuse of notation: as in the proof of [Lemma 3](#), we always use the same notation for an appraisal, its subalgebra extension, and its hyperallocation extension.

STEP 2: *Index the table entries, then introduce cutter appraisals and split rows.*

First, we index the table entries. In particular, for each  $j \in \{1, 2, \dots, s\}$ , we define (i)  $B'_j \equiv A_j \cap B'$ , (ii)  $B''_j \equiv A_j \cap (B \setminus B')$ , and (iii)  $B'''_j \equiv A_j \cap (C \setminus B)$ .

Second, let us say that a function  $\alpha : \mathcal{T} \rightarrow [0, 1]$  is a *cutter appraisal* if it satisfies the following requirements:

- Each row's value according to agent 1 is distributed across its entries: for each row  $j \in \{1, 2, \dots, s\}$ ,  $\alpha(B'_j) + \alpha(B''_j) + \alpha(B'''_j) = \alpha_1(A_j)$ .

- No sliver is assigned any value: for each  $S \in \mathcal{T} \cap \mathcal{S}^\circ$ ,  $\alpha(S) = 0$ .
- The ratio of the value in the first column to the value in the first two columns is precisely the target proportion:  $\sum_{j=1}^s \alpha(B'_j) = p_1 \cdot (\sum_{j=1}^s [\alpha(B'_j) + \alpha(B''_j)])$ .

Let  $\mathbb{A}$  denote the collection of cutter appraisals. Observe that a cutter appraisal need not be the restriction of a kitchen measure to  $\mathcal{T}$  because it may assign zero to a serving that is not a sliver.

Finally, for each  $\alpha \in \mathbb{A}$  and each  $j \in \{1, 2, \dots, s\}$ , we say that row  $j$  is  $\alpha$ -split if  $\alpha$  assigns positive value to at least two  $j$  entries:  $|\{B \in \{B'_j, B''_j, B'''_j\} \mid \alpha(B) > 0\}| > 1$ .

STEP 3: Construct  $\alpha_1^\mathcal{T} \in \mathbb{A}$  such that (i) there is at most one  $\alpha_1^\mathcal{T}$ -split row, and (ii) no value is assigned to the third column.

There are many ways to construct the desired  $\alpha_1^\mathcal{T}$ ; what follows is simply one concrete example. We begin with  $\mu_1|_{\mathcal{T}} \in \mathbb{A}$ , then apply our ROW POLARIZATION procedure to construct a cutter appraisal  $\alpha_1^\sim$  with at most one  $\alpha_1^\sim$ -split row, and finally modify  $\alpha_1^\sim$  to reach the desired  $\alpha_1^\mathcal{T}$ .

Formally, the ROW POLARIZATION procedure begins with  $\mu_1|_{\mathcal{T}} \in \mathbb{A}$  as the input to the first stage. At each stage, if the input  $\alpha \in \mathbb{A}$  is such that there is at most one  $\alpha$ -split row, then the procedure terminates with  $\alpha_1^\sim \equiv \alpha$ ; otherwise, it constructs the output  $\alpha' \in \mathbb{A}$  that serves as the input to the next stage as follows.

- Let  $j_1$  denote the smallest index  $\{1, 2, \dots, s\}$  of an  $\alpha$ -split row and let  $j_2$  denote the largest such index. Since there are at least two  $\alpha$ -split rows,  $j_1 \neq j_2$ .
- Define  $S_1^{\text{source}}$  and  $S_1^{\text{sink}}$  in row  $j_1$  as follows:
  - If  $\alpha(B'_{j_1}) > 0$  and  $\alpha(B''_{j_1}) > 0$ , then  $S_1^{\text{source}} \equiv B''_{j_1}$  and  $S_1^{\text{sink}} \equiv B'_{j_1}$ .
  - Else if  $\alpha(B'_{j_1}) > 0$  and  $\alpha(B'''_{j_1}) > 0$ , then  $S_1^{\text{source}} \equiv B'''_{j_1}$  and  $S_1^{\text{sink}} \equiv B'_{j_1}$ .
  - Else  $\alpha(B''_{j_1}) > 0$  and  $\alpha(B'''_{j_1}) > 0$ , and we define  $S_1^{\text{source}} \equiv B'''_{j_1}$  and  $S_1^{\text{sink}} \equiv B''_{j_1}$ .

Observe that transferring value from  $S_1^{\text{source}}$  to  $S_1^{\text{sink}}$  increases the ratio of the value in the first column to the value in the first two columns.

- Define  $S_2^{\text{source}}$  and  $S_2^{\text{sink}}$  in row  $j_2$  as follows:
  - If  $\alpha(B'_{j_2}) > 0$  and  $\alpha(B''_{j_2}) > 0$ , then  $S_2^{\text{source}} \equiv B'_{j_2}$  and  $S_2^{\text{sink}} \equiv B''_{j_2}$ .
  - Else if  $\alpha(B'_{j_2}) > 0$  and  $\alpha(B'''_{j_2}) > 0$ , then  $S_2^{\text{source}} \equiv B'_{j_2}$  and  $S_2^{\text{sink}} \equiv B'''_{j_2}$ .
  - Else  $\alpha(B''_{j_2}) > 0$  and  $\alpha(B'''_{j_2}) > 0$ , and we define  $S_2^{\text{source}} \equiv B'''_{j_2}$  and  $S_2^{\text{sink}} \equiv B''_{j_2}$ .

Observe that transferring value from  $S_2^{\text{source}}$  to  $S_2^{\text{sink}}$  decreases the ratio of the value in the first column to the value in the first two columns.

- For each value transfer rate  $\lambda \in \mathbb{R}_+$ , define  $\alpha_\lambda : \mathcal{T} \rightarrow [0, 1]$  as follows: (i) for each  $S \in \mathcal{T} \setminus \{S_1^{\text{source}}, S_1^{\text{sink}}, S_2^{\text{source}}, S_2^{\text{sink}}\}$ ,  $\alpha_\lambda(S) \equiv \alpha(S)$ , (ii)  $\alpha_\lambda(S_1^{\text{source}}) \equiv \alpha(S_1^{\text{source}}) - 1$ , (iii)  $\alpha_\lambda(S_1^{\text{sink}}) \equiv \alpha(S_1^{\text{sink}}) + 1$ , (iv)  $\alpha_\lambda(S_2^{\text{source}}) \equiv \alpha(S_1^{\text{source}}) - \lambda$ , and (v)  $\alpha_\lambda(S_2^{\text{sink}}) \equiv \alpha(S_1^{\text{sink}}) + \lambda$ .

- Let  $\lambda^*$  be the unique  $\lambda \in \mathbb{R}_+$  such that  $\sum_{j=1}^s \alpha_\lambda(B'_j) = p_1 \cdot (\sum_{j=1}^s [\alpha_\lambda(B'_j) + \alpha_\lambda(B''_j)])$ . It is easy to verify that  $\lambda^*$  is well-defined, that  $\lambda^*$  is greater than zero, and that  $\lambda^*$  can be explicitly constructed in cases given by the sinks and sources; we omit the details.
- Define  $v \equiv \min\{\alpha(S_1^{source}), \frac{1}{\lambda^*} \cdot \alpha(S_2^{source})\}$ , and define  $\alpha' : \mathcal{T} \rightarrow [0, 1]$  as follows: (i) for each  $S \in \mathcal{T} \setminus \{S_1^{source}, S_1^{sink}, S_2^{source}, S_2^{sink}\}$ ,  $\alpha'(S) \equiv \alpha(S)$ , (ii)  $\alpha'(S_1^{source}) \equiv \alpha(S_1^{source}) - v$ , (iii)  $\alpha'(S_1^{sink}) \equiv \alpha(S_1^{sink}) + v$ , (iv)  $\alpha'(S_2^{source}) \equiv \alpha(S_1^{source}) - \lambda^* \cdot v$ , and (v)  $\alpha'(S_2^{sink}) \equiv \alpha(S_1^{sink}) + \lambda^* \cdot v$ .

Intuitively,  $\lambda^*$  is the rate such that if we simultaneously and continuously transfer value (i) from source to sink in row  $j_1$  at a rate of one unit per second, and (ii) from source to sink in row  $j_2$  at a rate of  $\lambda^*$  units per second, then we preserve the ratio of the value in the first column to the value in the first two columns. We do so until one of the sources is drained of all value; necessarily  $\alpha'$  belongs to  $\mathbb{A}$ .

Observe that the ROW POLARIZATION procedure must terminate with  $\alpha_1^\approx$ . Indeed, at each stage for which the input and output differ, the output has more zero value assignments than the input, and the number of zero value assignments cannot exceed the total number of servings in the table.

Finally, we define  $\alpha_1^\mathcal{T}$  by modifying  $\alpha_1^\approx$  as follows: in each row  $j \in \{1, 2, \dots, s\}$ , transfer  $p \cdot \alpha(B_j''')$  from  $B_j'''$  to  $B_j'$  and transfer  $(1-p) \cdot \alpha(B_j''')$  from  $B_j'''$  to  $B_j''$ . It is easy to verify that the result has the desired properties.

STEP 4: *Conclude.*

First, we use  $\alpha_1^\mathcal{T}$  to construct an associated  $(\mathcal{P}^*, \alpha^*)$  that is  $\ell^+$ -deficient. If there are no  $\alpha_1^\mathcal{T}$ -split rows, then we simply take  $(\mathcal{P}^*, \alpha^*) \equiv (\mathcal{P}, \alpha)$ . Otherwise, (i) let  $j^*$  be the index of the unique  $\alpha_1^\mathcal{T}$ -split row, (ii) define  $p_1^* \in (0, 1)$  to be that row's ratio of the value in the first column to the value in the first two columns,  $p_1^* \equiv \frac{\alpha_1^\mathcal{T}(B_{j^*}')}{\alpha_1^\mathcal{T}(B_{j^*}') + \alpha_1^\mathcal{T}(B_{j^*}'')}$ , which for emphasis does not involve division by zero as  $j^*$  is  $\alpha_1^\mathcal{T}$ -split, (iii) let  $\mathbf{q}^* \equiv (1, \kappa, A_{j^*}, p_1^*)$ , and (iv) let  $(\mathcal{P}^*, \alpha^*)$  be the  $\ell^+$ -deficient member of  $\mathcal{R}(\mathbf{q}^* | \mathcal{P}, \alpha)$  promised by [Lemma 3](#) and constructed in its proof. In the latter case, there is  $A_{j^*}' \in \mathcal{P}^* \setminus \mathcal{P}$  such that

$$\begin{aligned}
\alpha_1^*(A_{j^*}') &= p_1^* \cdot \alpha_1^*(A_{j^*}) \\
&= \left[ \frac{\alpha_1^\mathcal{T}(B_{j^*}')}{\alpha_1^\mathcal{T}(B_{j^*}') + \alpha_1^\mathcal{T}(B_{j^*}'')} \right] \cdot [\alpha_1^\mathcal{T}(B_{j^*}') + \alpha_1^\mathcal{T}(B_{j^*}'') + \alpha_1^\mathcal{T}(B_{j^*}''')] \\
&= \left[ \frac{\alpha_1^\mathcal{T}(B_{j^*}')}{\alpha_1^\mathcal{T}(B_{j^*}') + \alpha_1^\mathcal{T}(B_{j^*}'')} \right] \cdot [\alpha_1^\mathcal{T}(B_{j^*}') + \alpha_1^\mathcal{T}(B_{j^*}'') + 0] \\
&= \alpha_1^\mathcal{T}(B_{j^*}').
\end{aligned}$$

Define  $A_{j^*}'' \equiv A_{j^*}' \setminus A_{j^*}'$ .

Second, we define  $f : \mathcal{P}^* \rightarrow \mathcal{T}$  such that for each  $S \in \mathcal{P}^*$ ,  $\alpha_1^*(S) = \alpha_1^\mathcal{T}(f(S))$ . First, for each  $j \in \{1, 2, \dots, s\}$  that is not  $\alpha_1^\mathcal{T}$ -split, (i)  $\alpha_1^*(A_j) = 0$  implies  $f(A_j)$  is the leftmost (as written here) member of  $\{B_j', B_j'', B_j'''\}$  that is not empty, and (ii)  $\alpha_1^*(A_j) > 0$  implies  $f(A_j)$  is the unique  $B \in \{B_j', B_j''\}$  such that  $\alpha_1^*(B) > 0$ . Second,  $A_{j^*}' \in \mathcal{P}^*$  implies  $f(A_{j^*}') \equiv B_{j^*}'$  and  $A_{j^*}'' \in \mathcal{P}^*$  implies  $f(A_{j^*}'') \equiv B_{j^*}''$ . It is straightforward to verify that  $f$  satisfies the desired property.

Third, we construct  $\alpha_2^\mathcal{T} : \mathcal{T} \rightarrow [0, 1]$  such that for each  $S \in \mathcal{P}^*$ ,  $\alpha_2^*(S) = \alpha_2^\mathcal{T}(f(S))$ . First, for each  $S \in \mathcal{P}^*$ , define  $\alpha_2^\mathcal{T}(f(S)) \equiv \alpha_2^*(S)$ . Second, for each  $S \in \mathcal{T} \setminus f(\mathcal{P}^*)$ , define  $\alpha_2^\mathcal{T}(S) \equiv 0$ .

For intuition, at this point we illustrate that if  $(\mathcal{T}, \alpha_1^\mathcal{T}, \alpha_2^\mathcal{T})$  belongs to  $\mathcal{R}(\mathbf{q}|\mathcal{P}, \alpha)$ , then we are done. Indeed, assume the hypothesis, let  $X \in \mathcal{X}_{\ell^+|\mathcal{T}}$ , and consider the associated hyperallocation  $X' \in \mathcal{X}_{\ell^+|\mathcal{P}^*}$  with  $r(X') = r(X)$  such that for each  $r \in \{1, 2, \dots, r(X')\}$ ,  $X'_r \equiv \cup\{f(S)|S \in \mathcal{P}^*, S \subseteq X_r\}$ . In other words,  $X'_r$  assigns each serving  $S$  to the agent to whom  $X_r$  assigns  $f(S)$ . Then (i) by  $\ell^+$ -deficiency of  $(\mathcal{P}^*, \alpha^*)$ ,  $\alpha_1^*(X) < e_1$  or  $\alpha_2^*(X) > e_1$ , and (ii) for each  $i \in N$ , we have  $\alpha_i^\mathcal{T}(X_r) = \alpha_i^*(X'_r)$ ; thus  $\alpha_1^\mathcal{T}(X_r) < e_1$  or  $\alpha_2^\mathcal{T}(X_r) > e_1$ . Since  $X \in \mathcal{X}_{\ell^+|\mathcal{T}}$  was arbitrary, we are done.

The proof is not yet complete because  $(\mathcal{T}, \alpha_1^\mathcal{T}, \alpha_2^\mathcal{T})$  need not belong to  $\mathcal{R}(\mathbf{q}|\mathcal{P}, \alpha)$ . Most importantly,  $\alpha_1^\mathcal{T}$  and  $\alpha_2^\mathcal{T}$  may violate sliver consistency by assigning zero to servings that are not slivers. To complete the proof, we modify  $(\mathcal{T}, \alpha_1^\mathcal{T}, \alpha_2^\mathcal{T})$  by moving a ‘tiny’ amount of value into these servings, then (if necessary) discarding any empty entries in  $\mathcal{T}$  so that we have a partition record. It is straightforward to verify that this can be done such that the result is an  $\ell^+$ -deficient ultraresponse, which completes the proof; the final three paragraphs provide an explicit construction for the interested reader.

To define tiny, we construct  $\varepsilon > 0$ . Since  $(\mathcal{P}^*, \alpha^*)$  is  $\ell^+$ -deficient, thus for each  $X \in \mathcal{X}_{\ell^+|\mathcal{P}^*}$ , either  $\alpha_1^*(X) < e_1$  or  $\alpha_2^*(X) > e_1$ , so the associated *deficit*

$$\delta(X) \equiv \min \left\{ \max\{0, e_1 - \alpha_1^*(X)\}, \max\{0, e_1 - \alpha_1^*(X)\} \right\}$$

is positive. First, define  $\varepsilon_\delta \equiv [\min_{X \in \mathcal{X}_{\ell^+|\mathcal{P}^*}} \delta(X)] \cdot \frac{1}{3s} \cdot \frac{1}{\ell^+} \cdot \frac{1}{2}$ , and observe that  $\varepsilon_\delta > 0$ . Intuitively, no shortage of an  $(\ell^+|\mathcal{P}^*)$ -hyperallocation can be overcome by giving each agent a bonus of  $\varepsilon_\delta$  for each of the  $3s$  table entries in each of  $\ell^+$  replicas. Second, define  $\varepsilon_\mathcal{T} \equiv \min[(\{\alpha_1(S)|S \in \mathcal{T}\} \cup \{\alpha_2(S)|S \in \mathcal{T}\}) \setminus \{0\}] \cdot \frac{1}{2} \cdot \frac{1}{2}$ . Intuitively, for each agent  $i$  and each table entry  $S$  that  $\alpha_i^\mathcal{T}$  assigns positive value, it is possible to transfer  $\varepsilon_\mathcal{T}$  from  $S$  to a second entry in that row, and then again to a third entry in that row, without draining all value from  $S$ . Finally, define  $\varepsilon \equiv \min\{\varepsilon_\delta, \varepsilon_\mathcal{T}\}$ . Observe that  $\varepsilon > 0$ .

Next, we use  $\varepsilon$  to construct  $(\mathcal{P}^+, \alpha^+) \in \mathcal{R}(\mathbf{q}|\mathcal{P}, \alpha)$  from  $(\mathcal{T}, \alpha^\mathcal{T})$ . First, for each  $i \in N$ , we construct  $\alpha_i^{\mathcal{T}^+} : \mathcal{T} \rightarrow [0, 1]$  from  $\alpha_i^\mathcal{T}$  by doing the following for each row  $j \in \{1, 2, \dots, s\}$ : (i) define  $\mathcal{S}_j^{\text{source}} \equiv \{B \in \{B'_j, B''_j, B'''_j\} | \alpha_i^\mathcal{T}(B) > 0\}$ , (ii) define  $\mathcal{S}_j^{\text{sink}} \equiv \{B \in \{B'_j, B''_j, B'''_j\} | B \notin \mathcal{S}_j^{\text{source}} \text{ and } B \notin \mathcal{S}^\circ\}$ , (iii) for each  $S \in \mathcal{S}_j^{\text{sink}}$ , define  $\alpha_i^{\mathcal{T}^+}(S) \equiv \varepsilon$ , (iv) for each  $S \in \mathcal{S}_j^{\text{source}}$ , define  $\alpha_i^{\mathcal{T}^+}(S) \equiv \alpha_i^\mathcal{T}(S) - \frac{|\mathcal{S}_j^{\text{sink}}| \cdot \varepsilon}{|\mathcal{S}_j^{\text{source}}|}$ , which for emphasis does not involve division by zero because we only call this definition when  $\mathcal{S}_j^{\text{source}}$  is nonempty, and (v) for each  $S \in \{B'_j, B''_j, B'''_j\}$  that does not belong to  $\mathcal{S}_j^{\text{source}} \cup \mathcal{S}_j^{\text{sink}}$ , define  $\alpha_i^{\mathcal{T}^+}(S) \equiv \alpha_i^\mathcal{T}(S)$ . Finally, we define  $\mathcal{P}^+ \equiv \mathcal{T} \setminus \{\emptyset\}$ , and for each  $i \in N$  we define  $\alpha_i^+ \equiv \alpha_i^{\mathcal{T}^+}|_{\mathcal{P}^+}$ .

Finally, we claim that  $(\mathcal{P}^+, \alpha^+) \in \mathcal{R}(\mathbf{q}|\mathcal{P}, \alpha)$  is  $\ell^+$ -deficient. First, by [Lemma 2](#) we have  $(\mathcal{P}^+, \alpha^+) \in \mathcal{R}(\mathbf{q}|\mathcal{P}, \alpha)$ . Second, let  $X \in \mathcal{X}_{\ell^+|\mathcal{P}^+}$ , and consider the associated hyperallocation  $X' \in \mathcal{X}_{\ell^+|\mathcal{P}^*}$  with  $r(X') = r(X)$  such that for each  $r \in \{1, 2, \dots, r(X')\}$ ,  $X'_r \equiv \cup\{f(S)|S \in \mathcal{P}^*, S \subseteq X_r\}$ . In other words,  $X'_r$  assigns each serving  $S$  to the agent to whom  $X_r$  assigns  $f(S)$ . Then (i) by definition of  $\varepsilon$ , there is  $i \in N$  such that  $3s\ell^+\varepsilon < e_i - \alpha_i^*(X')$ , and (ii) for each  $i \in N$ , we have  $\alpha_i^\mathcal{T}(X_r) = \alpha_i^*(X'_r)$ ; thus there is  $i \in N$  such that  $3s\ell^+\varepsilon < e_i - \alpha_i^\mathcal{T}(X)$ . Moreover, for each  $i \in N$  and each  $S \in \mathcal{T}$ ,  $\alpha_i^{\mathcal{T}^+}(S) \leq \alpha_i^\mathcal{T}(S) + \varepsilon$ ; thus as there are  $3s$  members of  $\mathcal{T}$  and  $r(X) \leq \ell^+$ , there is  $i \in N$  such that  $e_i - \alpha_i^{\mathcal{T}^+}(X) \geq e_i - \alpha_i^\mathcal{T}(X) - 3s\ell^+\varepsilon > 0$ . Finally,  $e_i - \alpha_i^+(X) = e_i - \alpha_i^{\mathcal{T}^+}(X) > 0$ . Since  $X \in \mathcal{X}_{\ell^+|\mathcal{T}}$  was arbitrary, we are done. ■

## Appendix 3

In this appendix, we prove [Proposition 1](#).

**PROPOSITION 1 (RESTATED):** Fix a setting. For each  $e \in E$  and each precision level  $\mathbf{p} \in \mathbb{N}$ ,  $P(e) > \mathbf{p}$  implies  $\text{cost}(e) > \lfloor \log_2 \log_2 2\mathbf{p} \rfloor$ .

**PROOF:** We first establish the proposition for two-agent settings, then conclude.

**STEP 1:** *Prove the proposition for two-agent settings.*

Fix a two-agent setting, let  $\mathbf{p} \in \mathbb{N}$ , and let  $e \in E$  be such that  $P(e) > \mathbf{p}$ . Since  $P(e) > \mathbf{p}$ , thus there is  $i \in N$  such that  $P_i(e) > \mathbf{p}$ ; denote the other agent by  $j$ . Finally, let  $\pi \in \Pi_e$ .

We begin by introducing some preliminary concepts. First, we say that each  $h \in H$  is *0-deficient*. Moreover, for each  $h \in H$  and each level  $\ell \in \mathbb{N}_0$ , we say that  $h$  is  *$\ell$ -deficient* if there is a valid  $\ell$ -deficient partition record  $(\mathcal{P}, \alpha)$  such that  $\mathbb{D}(h) \supseteq \mathbb{D}(\mathcal{P}, \alpha)$ . Second, define the initial level  $\ell_0 \equiv \mathbf{p}$ , and for each query count  $c \in \mathbb{N}_0$  define the level  $\ell_{c+1} \equiv \lfloor \sqrt{\ell_c/2} \rfloor$ . Finally, for each  $c \in \mathbb{N}_0$  and each  $\mu \in \mathbb{D}$ , define  $h_c(\mu) \in H \cup \{\text{error}\}$  as follows: (i) if  $c \leq \text{cost}(\mu)$ , then  $h_c(\mu)$  is the history in the play determined by  $\mu$  and  $\pi$  whose query count is  $c$ , and (ii) otherwise  $h_c(\mu) \equiv \text{error}$ .

We claim that for each  $c \in \mathbb{N}_0$  such that  $\ell_c \geq 1$ , there is  $\mu \in \mathbb{D}$  such that for each  $c' \in \{0, 1, \dots, c+1\}$ ,  $h_{c'}(\mu) \in H$  and  $h_{c'}(\mu)$  is  $\ell_{c'}$ -deficient. We proceed by induction. For the base step, since  $P_i(e) > \mathbf{p}$ , thus  $e_i$  cannot be written as a fraction whose denominator is at most  $\ell_0 = \mathbf{p}$ , so the unique valid partition record for partition  $\{C\}$  is  $\ell_0$ -deficient, so for each  $\mu \in \mathbb{D}$  we have that  $h_0(\mu)$  is  $\ell_0$ -deficient; thus the claim holds for  $c = 0$ . For the inductive step, let  $c \in \mathbb{N}_0$  be such that the claim holds with measure profile  $\mu$ . If  $\ell_{c+1} < 1$  then we are done, so assume  $\ell_{c+1} \geq 1$ . By the inductive hypothesis,  $h_{c+1}(\mu)$  is  $\ell_{c+1}$ -deficient, so there is a valid  $\ell_{c+1}$ -deficient partition record  $(\mathcal{P}, \alpha)$  such that  $\mathbb{D}(h_{c+1}(\mu)) \supseteq \mathbb{D}(\mathcal{P}, \alpha)$ . Since  $\ell_{c+1} \geq 1$ , there is no  $X \in \mathcal{X}$  such that for each  $\mu \in \mathbb{D}(\mathcal{P}, \alpha)$ ,  $X$  is  $(e|\mu)$ -proportional; thus since  $\pi \in \Pi_e$ , necessarily  $\pi(h_{c+1}(\mu))$  is a query  $\mathbf{q}$ . By [Lemma 4](#), there is  $\mu' \in \mathbb{D}(\mathcal{P}, \alpha)$  such that the ultrareponse  $r(\mathbf{q}|\mu', \mathcal{P}, \alpha)$  is  $\lfloor \sqrt{\ell_{c+1}/2} \rfloor$ -deficient, or equivalently  $\ell_{c+2}$ -deficient. Since  $\mu' \in \mathbb{D}(\mathcal{P}, \alpha) \subseteq \mathbb{D}(h_{c+1}(\mu))$ , thus  $\mu$  and  $\mu'$  share the same  $c+1$  queries and responses. Altogether, then, the claim holds for  $c+1$  with measure profile  $\mu'$ , so as  $c \in \mathbb{N}_0$  was arbitrary we are done.

To conclude, define  $c^* \equiv \max\{c \in \mathbb{N}_0 \mid \mathbf{p} \geq 2^{2^{c-1}}\}$ , which is well-defined since  $\mathbf{p} \geq 1$ . We claim that for each  $c \in \{0, 1, \dots, c^*\}$ ,  $\ell_c \geq 2^{2^{c^*-c}-1}$ . We proceed by induction. For the base step,  $\ell_0 = \mathbf{p} \geq 2^{2^{c^*-0}-1}$ , as desired. For the inductive step, let  $c \in \{0, 1, \dots, c^*-1\}$  be such that  $\ell_c \geq 2^{2^{c^*-c}-1}$ . Then  $\ell_{c+1} = \lfloor \sqrt{\ell_c/2} \rfloor \geq \lfloor \sqrt{2^{2^{c^*-c}-1}/2} \rfloor = \lfloor \sqrt{2^{2^{c^*-c}-2}} \rfloor = \lfloor 2^{2^{c^*-c}-1} \rfloor = 2^{2^{c^*-c}-1}$ , so the claim holds for  $c+1$ . Since  $c \in \{0, 1, \dots, c^*-1\}$  was arbitrary, we are done. In particular, we have that  $\ell_{c^*} \geq 1$ , so by the previous paragraph there is  $\mu \in \mathbb{D}$  such that  $h_{c^*+1}(\mu) \in H$ , so  $\text{cost}(\pi|\mu) > c^*$ , so  $\text{cost}(\pi) > c^*$ . Since  $\pi \in \Pi_e$  was arbitrary, thus  $\text{cost}(e) > c^*$ . Finally, by definition of  $c^*$  we have that  $c^* \geq \lfloor \log_2 \log_2 2\mathbf{p} \rfloor$ . Altogether, then, we have  $\text{cost}(e) > \lfloor \log_2 \log_2 2\mathbf{p} \rfloor$ , as desired.

**STEP 2:** *Conclude.*

Assume, by way of contradiction, there is a setting  $(C, \mathbb{K}, N)$  for which the proposition

does not hold: there are  $\mathbf{p} \in \mathbb{N}$ ,  $e \in E$  with  $P(e) > \mathbf{p}$ , and  $\pi \in \Pi_e$  such that for each  $\mu \in \mathbb{D}$ ,  $\text{cost}(\pi|\mu) \leq \lfloor \log_2 \log_2 2\mathbf{p} \rfloor$ . For each setting with one agent, each entitlement profile's precision index is 1 and thus the proposition holds vacuously, so  $n \neq 1$ . Moreover, for each setting with two agents, the proposition holds by Step 1, so  $n \neq 2$ . Altogether, then,  $n > 2$ . We refer to  $(C, \mathbb{K}, N)$  as the *large setting*. Since  $P(e) > \mathbf{p}$ , thus there is  $i \in N$  such that  $P_i(e) > \mathbf{p}$ . Select  $j \in N \setminus \{i\}$  and define  $N' \equiv \{i, j\}$ . We refer to  $(C, \mathbb{K}, N')$  with  $|N'| = 2$  as the *small setting*.

Since our notation suppresses the setting while the following argument involves two settings, we simply use  $'$  to distinguish objects for the small setting from the associated objects from the large setting—for example, writing  $E$  for the large setting and  $E'$  for the small setting.

To complete this step, we prove that the proposition does not hold for the small setting, which contradicts Step 1. More precisely, define  $e' \in E'$  by  $(e'_i, e'_j) \equiv (e_i, 1 - e_i)$ . By construction,  $P'(e') = P(e) > \mathbf{p}$ . We claim that  $\text{cost}'(e') \leq \lfloor \log_2 \log_2 2\mathbf{p} \rfloor$ , which we establish by exhibiting a protocol  $\pi' \in \Pi'_{e'}$  such that  $\text{cost}(\pi') \leq \lfloor \log_2 \log_2 2\mathbf{p} \rfloor$ . Intuitively, the protocol  $\pi'$  that we exhibit imitates  $\pi$  in a suitable sense.

We begin by introducing some objects that we use to construct  $\pi'$ . First, select a *guess function*  $\mathbf{g} : 2^{\mathbb{D}'} \setminus \{\emptyset\} \rightarrow \mathbb{D}'$  that selects from each nonempty domain  $\mathbb{D}^* \subseteq \mathbb{D}'$  a guess  $\mathbf{g}(\mathbb{D}^*) \in \mathbb{D}^*$ . Second, define  $\mathbb{D}^i \equiv \{\mu \in \mathbb{D} \mid \text{for each pair } k, k' \in N \setminus \{i\}, \mu_k = \mu_{k'}\}$ , and define the *extension function*  $\mathcal{E} : \mathbb{D}' \rightarrow \mathbb{D}^i$  be such that for each  $\mu' \in \mathbb{D}'$ ,  $\mathcal{E}(\mu')$  is the  $\mu \in \mathbb{D}^i$  such that (i)  $\mu_i = \mu'_i$ , and (ii) for each  $k \in N \setminus \{i\}$ ,  $\mu_k = \mu'_j$ . Finally, for each  $c \in \mathbb{N}_0$  and each  $\mu \in \mathbb{D}$ , we define  $h_c(\pi|\mu) \in H \cup \{\text{error}\}$  as follows: (i) if  $c \leq \text{cost}(\pi|\mu)$ , then  $h_c(\pi|\mu)$  is the history in the play determined by  $\mu$  and  $\pi$  whose query count is  $c$ , and (ii) otherwise  $h_c(\pi|\mu) \equiv \text{error}$ .

Next, we construct  $\pi'$ . In particular, for each  $h' \in H'$ , (i) let  $c$  denote the query count at  $h'$ , (ii) define the guess  $\mathbf{g}' \equiv \mathbf{g}(\mathbb{D}'(h'))$ , (iii) define the guess extension  $\mathbf{g} \equiv \mathcal{E}(\mathbf{g}')$ , and (iv) define  $\pi'(h') \in \mathcal{Q}' \cup \mathcal{X}'$  in three cases as follows.

- If  $c < \text{cost}(\pi|\mathbf{g})$ , then  $\pi(h_c(\pi|\mathbf{g})) \in \mathcal{Q}$ . If the cutter of this query belongs to  $N'$ , then define  $\pi'(h')$  to be the same query; otherwise define  $\pi'(h')$  to be query formed by taking this query and then making  $j$  the cutter.
- If  $c = \text{cost}(\pi|\mathbf{g})$ , then define  $X \equiv \text{out}(\pi|\mathbf{g})$ , let  $\pi'(h') \in \mathcal{X}'$  be defined by  $X'_i = X_i$  and  $X'_j \equiv C \setminus X'_i$ , and define  $\pi'(h') \equiv X'$ . Since  $\pi \in \Pi_e$ , thus  $X$  is  $(e|\mathbf{g})$ -proportional, from which it follows that  $X'$  is  $(e'|\mathbf{g}')$ -proportional.
- If  $c > \text{cost}(\pi|\mathbf{g})$ , then  $h'$  will ultimately be off-path. For concreteness, set  $\pi'(h') \in \mathcal{X}'$  to be the allocation at which agent  $i$  receives the entire cake.

Observe that since  $\pi'$  assigns the same action to any two histories with the same chronicle, as each history's guess only depends on its chronicle, thus  $\pi'$  is indeed a protocol. For each  $c \in \mathbb{N}_0$  and each  $\mu' \in \mathbb{D}'$ , we define  $h'_c(\pi'|\mu') \in H' \cup \{\text{error}\}$  analogously to the associated notation for  $\pi$ .

*Claim.* Fix an arbitrary *true* measure profile  $\mathbf{t}' \in \mathbb{D}'$ . Let us write  $\mathbf{t} \equiv \mathcal{E}(\mathbf{t}')$ , and for each  $c \in \mathbb{N}_0$  let us write  $h'_c \equiv h'_c(\pi'|\mathbf{t}')$ ,  $\mathbf{g}'_c \equiv \mathbf{g}(\mathbb{D}'(h'_c))$ ,  $\mathbf{g}_c \equiv \mathcal{E}(\mathbf{g}'_c)$ ,  $h_c^t \equiv h_c(\pi|\mathbf{t})$ , and  $h_c^g \equiv h_c(\pi|\mathbf{g}_c)$ . We claim that for each  $c \in \mathbb{N}_0$  such that  $h'_c \in H'$ , we have

- $h_c^t \in H$  and  $h_c^g \in H$ ,
- $\mathbb{D}(h_c^t) = \mathbb{D}(h_c^g)$ , and

- $\mathbb{D}(h_c^t) \cap \mathbb{D}^i = \mathcal{E}(\mathbb{D}'(h'_c))$ .

*Proof of Claim.* We prove the claim by induction on the query count, and the base step  $c = 0$  is trivial; thus let us make the inductive hypothesis that the claim holds for  $c \in \mathbb{N}_0$ . If  $h'_{c+1} = \text{error}$  then we are done; thus let us assume  $h'_{c+1} \in H'$ , in which case  $\pi'(h'_c) \in \mathcal{Q}'$ . Then by construction of  $\pi'$ ,  $h_c^g \in H$  and  $\pi(h_c^g) \in \mathcal{Q}$ , so  $h_{c+1}^g \in H$ . Moreover, (i) since  $t'$  and  $g'_c$  both belong to  $\mathbb{D}'(h'_c)$ , thus by the inductive hypothesis  $t$  and  $g_c$  both belong to  $\mathbb{D}(h_c^t)$ , and (ii)  $g_c$  belongs to  $\mathbb{D}(h_c^g)$ ; thus necessarily  $\mathbb{D}(h_c^t) = \mathbb{D}(h_c^g)$ .<sup>15</sup> Altogether, then,  $\pi(h_c^t) = \pi(h_c^g) \in \mathcal{Q}$ , so  $h_{c+1}^t \in H$ . This establishes the first conclusion, from which the second conclusion follows by definition of  $g_{c+1}$ . To conclude, define  $q \equiv \pi(h_c^t) = \pi(h_c^g)$  and define  $q' \equiv \pi'(h'_c)$ . Using the definition of  $\pi'$ , it is straightforward to verify that whether or not the cutter of this query belongs to  $N'$ , we have  $\mathcal{E}(r'(q'|t')) = r(q|t) \cap \mathbb{D}^i$ , which together with the inductive hypothesis directly implies the third conclusion. Altogether, then, the claim holds for  $c + 1$ , so as  $c \in \mathbb{N}_0$  was arbitrary we are done.

*Proof from Claim.* To conclude, let  $t' \in \mathbb{D}'$ . Define  $t \equiv \mathcal{E}(t')$ , define  $X' \equiv \text{out}(\pi'|t')$  and define  $X \equiv \text{out}(\pi|t)$ . We claim that  $\text{cost}'(\pi'|t') \leq \lfloor \log_2 \log_2 2p \rfloor$  and  $X'$  is  $(e'|t')$ -proportional. Indeed, by our hypotheses about  $\pi$ , we have that  $\text{cost}(\pi|t) \leq \lfloor \log_2 \log_2 2p \rfloor$ ,  $X \in \mathcal{X}$ , and  $X$  is  $(e|t)$ -proportional. By our Claim and the definition of  $\pi'$ , the sequence of actions in the play determined by  $\pi'$  and  $t'$  is formed from the sequence of actions in the play determined by  $\pi$  and  $t$  by (i) taking any query whose cutter does not belong to  $N'$  and making its cutter  $j$ , and (ii) replacing  $X$  with the allocation  $X' \in \mathcal{X}$  such that  $X'_i = X_i$ . Then  $\text{cost}'(\pi'|t') = \text{cost}(\pi|t) \leq \lfloor \log_2 \log_2 2p \rfloor$ . Moreover, since  $X$  is  $(e|t)$ -proportional, thus  $t'_i(X_i) \geq e_i = e'_i$ ; it directly follows that  $X'$  is  $(e'|t')$ -proportional. Since  $t' \in \mathbb{D}'$  was arbitrary, thus  $\pi' \in \Pi'_{e'}$  and  $\text{cost}'(\pi') \leq \lfloor \log_2 \log_2 2p \rfloor$ , so  $\text{cost}'(e') \leq \lfloor \log_2 \log_2 2p \rfloor$ . But then the small setting is a two-agent setting for which the proposition does not hold, contradicting Step 1. ■

## Appendix 4

In this appendix, we prove [Theorem 2](#).

To begin, the following lemma gathers two slight extensions of exercises in [Savage \(1954\)](#); the proof is omitted.

LEMMA 5: Fix a kitchen and a qualitative probability. We have the following.

- For each  $A \in \mathcal{S}$  and each pair  $B, B' \in \mathcal{S}$  such that  $B, B' \subseteq A$ , we have  $B \succsim B'$  if and only if  $(A \setminus B') \succsim (A \setminus B)$ .
- For each four  $A, A', B, B' \in \mathcal{S}$ , if (i)  $A \cap A' = \emptyset$ , (ii)  $A \succsim B$ , and (iii)  $A' \succsim B'$ , then  $(A \cup A') \succsim (B \cup B')$ . If moreover  $A \succ B$ , then  $(A \cup A') \succ (B \cup B')$ .

We first establish that under the classic qualitative probability axioms, knife continuity is stronger than tightness.

<sup>15</sup>Indeed, this follows from two observations: (i) if the domains of two histories intersect, then they are nested, and (ii) if the domains of two histories are strictly nested, then their query counts are not the same.

PROPOSITION 3: Fix a kitchen. If a qualitative probability satisfies *knife continuity*, then it satisfies *tightness*. Moreover, for each  $\kappa \in \mathbb{K}$  and each pair  $C^*, A \in \mathcal{S}$  such that  $C^* \succsim A$ , there is  $t \in [0, 1]$  such that  $(C^* \cap \kappa_t) \sim A$ .

PROOF: Let  $\succsim$  satisfy the hypotheses and let  $\kappa \in \mathbb{K}$ .

STEP 1: For each  $A \in \mathcal{S}$ , there is  $t \in [0, 1]$  such that  $\kappa_t \sim A$ .

Let  $A \in \mathcal{S}$ , define  $T^+ \equiv \{t \in [0, 1] \mid \kappa_t \succsim A\}$ , and define  $T^- \equiv \{t \in [0, 1] \mid A \succsim \kappa_t\}$ . By *knife continuity*,  $T^+$  and  $T^-$  are closed, and by *completeness*  $T^+ \cup T^- = [0, 1]$ , so since  $[0, 1]$  is connected there is  $t \in T^+ \cap T^-$ , and necessarily  $\kappa_t \sim A$ , as desired.

STEP 2: For each pair  $C^*, A \in \mathcal{S}$  such that  $C^* \succsim A$ , there is  $t \in [0, 1]$  such that  $(C^* \cap \kappa_t) \sim A$ .

Let  $C^*, A \in \mathcal{S}$  such that  $C^* \succsim A$ . To ease notation, for each  $x \in [0, 1]$ , we define  $\kappa_x^* \equiv (C^* \cap \kappa_x)$ . We use  $x$  to index cuts of  $C^*$  and  $y$  to index cuts of  $C$ ; thus we write  $\{\kappa_x^*\}_{x \in [0, 1]}$  and  $\{\kappa_y\}_{y \in [0, 1]}$ .

To begin, define  $x^\leftarrow \equiv \inf\{x \in [0, 1] \mid \kappa_x^* \succsim A\}$  and  $x^\rightarrow \equiv \sup\{x \in [0, 1] \mid A \succsim \kappa_x^*\}$ ; by *knife continuity*, *monotonicity*, and *separability*, both  $x^\leftarrow$  and  $x^\rightarrow$  are well-defined. By *completeness*, we have  $x^\rightarrow \geq x^\leftarrow$ , and if  $x^\rightarrow > x^\leftarrow$  then for  $x \equiv \frac{x^\leftarrow + x^\rightarrow}{2}$  we have  $\kappa_x^* \sim A$  and we are done; thus let us assume  $x^\rightarrow = x^\leftarrow$ . Define  $x^\circ \equiv x^\rightarrow = x^\leftarrow$ , define  $\kappa_\circ^* \equiv \kappa_{x^\circ}^*$ , and define  $\kappa_\circ \equiv \kappa_{x^\circ}$ .

First, we claim that  $A \succsim \kappa_\circ^*$ . Indeed, assume by way of contradiction that  $\kappa_\circ^* \succ A$  and define  $B \equiv (\kappa_\circ \setminus C^*)$ . Since  $\kappa_\circ^* = (\kappa_\circ \cap C^*) \subseteq (\kappa_\circ \cap C^*) \cup (C \setminus \kappa_\circ) = (C \setminus B)$ , thus by *monotonicity* we have  $(C \setminus B) \succsim \kappa_\circ^* \succ A$ . By Step 1, there are  $y(A), y(C \setminus B) \in [0, 1]$  such that  $\kappa_{y(A)} \sim A$  and  $\kappa_{y(C \setminus B)} \sim (C \setminus B)$ , and moreover by *monotonicity* we have  $y(C \setminus B) > y(A)$ . Then  $\kappa_{y(A)}$  and  $(C \setminus \kappa_{y(C \setminus B)})$  are disjoint, and by Lemma 5 the former is equivalent to  $A$  and the latter is equivalent to  $B$ ; define  $A \oplus B$  to be their union  $\kappa_{y(A)} \cup (C \setminus \kappa_{y(C \setminus B)})$ . For each  $x \in [0, x^\circ)$ , we have  $x < x^\leftarrow$  and thus  $\kappa_\circ^* \succ A \succ \kappa_x^*$ , so by Lemma 5 we have  $\kappa_\circ = \kappa_\circ^* \cup B \succ A \oplus B \succ \kappa_x^* \cup B \succsim \kappa_x$ . But then by *monotonicity* we have that  $\{t \in [0, 1] \mid A \oplus B \succsim \kappa_t\} = [0, x^\circ)$ , so this set is not closed, contradicting *knife continuity*.

To conclude, we first claim that  $\kappa_\circ^* \succsim A$ . Indeed, the argument is symmetric: we omit the details, but remark that if for each  $t \in [0, 1]$  we define  $\kappa'_t \equiv (C \setminus \kappa_{1-t})$ , then we can use the same argument using  $\kappa'$  instead of  $\kappa$ . Altogether, then,  $C^* \cap \kappa_{x^\circ} = \kappa_\circ^* \sim A$ , as desired.

STEP 3: Conclude.

We claim that for each pair  $A, B \in \mathcal{S}$ , if for each  $A' \in \mathcal{S}$  such that  $A' \succ \emptyset$  and  $A \cap A' = \emptyset$  we have  $A \cup A' \succsim B$ , then  $A \succsim B$ . Indeed, let  $A$  and  $B$  satisfy the hypothesis. By Step 1, there are  $t, t' \in [0, 1]$  such that  $\kappa_t \sim A$  and  $\kappa_{t'} \sim B$ ; thus we can define  $t_A \equiv \sup\{t \in [0, 1] \mid A \sim \kappa_t\}$  and  $t_B \equiv \inf\{t \in [0, 1] \mid B \sim \kappa_t\}$ . By *knife continuity*,  $\kappa_{t_A} \sim A$  and  $\kappa_{t_B} \sim B$ . Assume, by way of contradiction, that  $t_B > t_A$ , and define  $t^* \equiv \frac{t_A + t_B}{2}$ . By definition of  $t_A$  and  $t_B$  and *monotonicity*, we have  $B \succ \kappa_{t^*} \succ A \sim \kappa_{t_A}$ ; thus by *separability* and *monotonicity* we have  $(C \setminus A) \succsim (\kappa_{t^*} \setminus \kappa_{t_A}) \succ \emptyset$ . By Step 2, there is  $t \in [0, 1]$  such that for  $A' \equiv (C \setminus A) \cap \kappa_t$ , we have  $A' \sim (\kappa_{t^*} \setminus \kappa_{t_A}) \succ \emptyset$ . By Lemma 5 we

have  $A \cup A' \sim \kappa_{t_A} \cup (\kappa_{t^*} \setminus \kappa_{t_A}) = \kappa_{t^*}$ , and by hypothesis we have  $A \cup A' \succsim B$ . But then  $\kappa_{t^*} \succsim B$ , contradicting  $B \succ \kappa_{t^*}$ . Altogether, then,  $t_A \geq t_B$ , so by *monotonicity* we have  $A \sim \kappa_{t_A} \succsim \kappa_{t_B} \sim B$ , as desired.

To conclude, *tightness* follows directly from the claim. ■

The theorem follows easily from the proposition.

**THEOREM 2 (RESTATED):** Fix a kitchen. A qualitative probability  $\succsim$  satisfies *fineness*, *sliver nullity*, and *knife continuity* if and only if there is a kitchen measure  $\mu_0$  such that for each pair  $A, B \in \mathcal{S}$ , we have  $A \succsim B$  if and only if  $\mu_0(A) \geq \mu_0(B)$ . In this case,  $\mu_0$  is the unique such kitchen measure.

**PROOF:** It is straightforward to verify that if  $\succsim$  has a kitchen measure representation, then it satisfies the axioms; we omit the argument. Thus let us assume that  $\succsim$  satisfies *fineness*, *sliver nullity*, and *knife continuity*. By [Proposition 3](#),  $\succsim$  satisfies *tightness*, so by [Theorem WM](#) it has a unique probability measure representation  $\mu_0$  whose range  $R \equiv \{\mu_0(A) | A \in \mathcal{S}\}$  is a dense subset of the unit interval. Let  $\kappa \in \mathbb{K}$  and let  $f : [0, 1] \rightarrow [0, 1]$  be such that for each  $t \in [0, 1]$ ,  $f(t) = \mu_0(\kappa_t)$ . For each  $v \in R$ , there is  $A_v \in \mathcal{S}$  such that  $\mu_0(A_v) = v$ , so by [Proposition 3](#) there is  $t \in [0, 1]$  such that  $\kappa_t = (C \cap \kappa_t) \sim A_v$  and thus  $\mu_0(\kappa_t) = v$ ; thus  $v \in f([0, 1])$ . Since  $v \in R$  was arbitrary, thus  $R \subseteq f([0, 1])$ , so  $R = f([0, 1])$ . Then the range of  $f$  is a dense subset of the unit interval, and moreover by *monotonicity*  $f$  is non-decreasing, so by a classic result<sup>16</sup> we have that  $f$  is continuous. Since  $f(0) = 0$  and  $f(1) = 1$ , thus  $0 \in R$  and  $1 \in R$ , so by the Intermediate Value Theorem we have that  $R = f([0, 1]) = [0, 1]$ . From here, it follows directly from *sliver nullity* and [Proposition 3](#) that  $\mu_0$  is a kitchen measure. ■

## Appendix 5

In this appendix, we prove [Theorem 3](#).

**THEOREM 3 (RESTATED):** Fix a setting and let  $e \in E$ . If all entitlements are rational, then the game has a value and this value is achieved in a Nash equilibrium. If there is an irrational entitlement, then the game's value is  $-\infty$  but the game has no Nash equilibrium.

**PROOF:** Fix a setting and let  $e \in E$ .

First, we associate each protocol with a mediator strategy in the adversary game. Indeed, select an arbitrary allocation  $X \in \mathcal{X}$ . For each protocol  $\pi \in \Pi$ , let  $\sigma_\pi \in \Sigma_m$  denote the following strategy: at each history  $h$  in the adversary game (i) if  $\mathbb{D}(h) \neq \emptyset$ , then let  $\sigma_\pi(h)$  denote the action taken at the information set in the division game whose chronicle is  $h$ , and (ii) if  $\mathbb{D}(h) = \emptyset$ , then  $\sigma_\pi(h) = X$ .

Second, we associate each mediator strategy in the adversary game with a protocol. Indeed, for each  $\sigma \in \Sigma_m$ , let  $\pi_\sigma \in \Pi$  denote the following protocol: at each information set in the division game, let  $h$  denote the associated chronicle and select the action  $\sigma(h)$ .

<sup>16</sup>In particular, if  $f : [0, 1] \rightarrow [0, 1]$  has dense range and is non-decreasing, then it is continuous. This result is sometimes called Froda's Theorem; see Theorem 4.30 and its corollary in [Rudin \(1976\)](#).

Finally, we associate each measure profile with an adversary strategy. Indeed, for each measure profile  $\mu \in \mathbb{D}$ , let  $\sigma_\mu \in \Sigma_a$  denote the following strategy: after each query  $\mathbf{q}$ , select  $\mathbf{r}(\mathbf{q}|\mu)$ .

CASE 1: *All entitlements are rational.* Define  $c \equiv \text{cost}(e)$ . By definition, there is an optimal protocol  $\pi^*$  such that  $\text{cost}(\pi) = c$ . Moreover, since there is a bounded protocol (Steinhaus, 1948), thus  $c \in \mathbb{N}_0$ , so there is  $\mu^* \in \mathbb{D}$  such that  $\text{cost}(\pi^*|\mu^*) = c$ . It is straightforward to verify that (i)  $(\sigma_{\pi^*}, \sigma_{\mu^*})$  is a Nash equilibrium, (ii)  $\mathbb{P}(\sigma_{\pi^*}, \sigma_{\mu^*}) = c$ , and (iii)  $c$  is the value of the adversary game.

CASE 2: *There is an irrational entitlement.* First, we claim that the adversary game has no Nash equilibrium. Indeed, assume by way of contradiction that  $(\sigma_m, \sigma_a)$  is a Nash equilibrium. We must have  $\mathbb{P}(\sigma_m, \sigma_a) = -\infty$ , else since  $\pi_{\sigma_m}$  is not bounded (Corollary 1), there is  $\mu \in \mathbb{D}$  such that  $\text{cost}(\pi_{\sigma_m}|\mu) > -\mathbb{P}(\sigma_m, \sigma_a)$ , so  $\mathbb{P}(\sigma_m, \sigma_\mu) = -\text{cost}(\pi_{\sigma_m}|\mu) < \mathbb{P}(\sigma_m, \sigma_a)$ , contradicting that  $(\sigma_m, \sigma_a)$  is a Nash equilibrium. But then since there is a finite protocol  $\pi$  (Barbanel, 1995), necessarily  $\mathbb{P}(\sigma_\pi, \sigma_a) > -\infty = \mathbb{P}(\sigma_m, \sigma)$ , contradicting that  $(\sigma_m, \sigma_a)$  is a Nash equilibrium.

Second, we claim that  $\sup_{\sigma_m \in \Sigma_m} \inf_{\sigma_a \in \Sigma_a} \mathbb{P}(\sigma_m, \sigma_a) = -\infty$ . Indeed, for each  $\sigma_m \in \Sigma_m$  and each  $c \in \mathbb{N}_0$ , since  $\pi_{\sigma_m}$  is not bounded (Corollary 1) there is  $\mu \in \mathbb{D}$  such that  $-\mathbb{P}(\sigma_m, \sigma_\mu) = \text{cost}(\pi_{\sigma_m}|\mu) > c$ , so  $\inf_{\sigma_a \in \Sigma_a} \mathbb{P}(\sigma_m, \sigma_a) \leq \mathbb{P}(\sigma_m, \sigma_\mu) < -c$ . Since  $c \in \mathbb{N}_0$  was arbitrary, thus  $\inf_{\sigma_a \in \Sigma_a} \mathbb{P}(\sigma_m, \sigma_a) = -\infty$ . Since  $\sigma_m \in \Sigma_m$  was arbitrary, we are done.

Finally, we claim that  $\inf_{\sigma_a \in \Sigma_a} \sup_{\sigma_m \in \Sigma_m} \mathbb{P}(\sigma_m, \sigma_a) = -\infty$ . Indeed, let  $c^* \in \mathbb{N}_0$  denote an arbitrary cost; we exhibit an adversary strategy  $\sigma_{c^*}$  that guarantees a higher cost. First, define  $\ell_0 \equiv 2^{3^{c^*}}$ , and for each  $c \in \{0, 1, \dots, c^* - 1\}$  define  $\ell_{c+1} \equiv \lfloor \sqrt{\ell_c/2} \rfloor$ . It is straightforward to verify by induction that for each  $c \in \{0, 1, \dots, c^*\}$ , we have  $\ell_c \geq 2^{2^{c^*-c}}$ , and in particular  $\ell_{c^*} \geq 2$ . We construct  $\sigma_{c^*}$  sequentially: at each adversary history  $h$  such that we have already defined the adversary's actions at all earlier adversary histories, define  $\sigma_{c^*}(h)$  as follows.

- If (i)  $\mathbb{D}(h) \neq \emptyset$ , (ii) there have thus far been  $c \in \{0, 1, \dots, c^* - 1\}$  queries, and (iii) there is an  $\ell_c$ -deficient partition record  $(\mathcal{P}, \alpha)$  that extends the record given by all previous responses, then if  $\mathbf{q}$  denotes the most recent (and only unanswered) query, Lemma 4 promises there is  $(\mathcal{P}^+, \alpha^+) \in \mathcal{R}(\mathbf{q}|\mathcal{P}, \alpha)$  that is  $\ell_{c+1}$ -deficient. Select a measure  $\mu \in \mathbb{D}$  associated with this ultrasponse, then select the response  $\mathbf{r}(\mathbf{q}|\mu)$ .
- Otherwise, select any response.

This completes our definition of  $\sigma_{c^*}$ . By construction, for each mediator strategy  $\sigma_m \in \Sigma_m$  and each  $c \in \{0, 1, \dots, c^*\}$ , if there is a mediator history  $h$  at which  $c$  queries have been already asked that is on-path given  $(\sigma_m, \sigma_{c^*})$ , then the record of all previous responses can be extended to a partition record  $(\mathcal{P}, \alpha)$  that is  $\ell_c$ -deficient, and moreover  $\ell_c \geq \ell_{c^*} \geq 2 > 1$ , so there is no allocation that is  $e$ -proportional for each member of  $\mathbb{D}(\mathcal{P}, \alpha)$  and  $\mathbb{D}(\mathcal{P}, \alpha) \subseteq \mathbb{D}(h)$ . Since  $c \in \{0, 1, \dots, c^*\}$  was arbitrary, thus  $\mathbb{P}(\sigma_m, \sigma_{c^*}) < -c^*$ . Since  $\sigma_m \in \Sigma_m$  was arbitrary, thus  $\inf_{\sigma_a \in \Sigma_a} \sup_{\sigma_m \in \Sigma_m} \mathbb{P}(\sigma_m, \sigma_a) \leq \sup_{\sigma_m \in \Sigma_m} \mathbb{P}(\sigma_m, \sigma_{c^*}) < -c^*$ . Since  $c^* \in \mathbb{N}_0$  was arbitrary, we are done. ■

# References

- AKBARPOUR, MOHAMMAD AND SHENGWU LI (2020). “Credible auctions: a trilemma.” *Econometrica* 88, 425-467.
- AZIZ, HARIS AND SIMON MACKENZIE (2016). “A discrete and bounded envy-free cake cutting protocol for any number of agents.” In: *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, 416-427.
- BARBANEL, JULIUS (1995). “Game-theoretic algorithms for fair and strongly fair cake division with entitlements.” *Colloquium Mathematicum* 69, 59-73.
- BARBANEL, JULIUS AND ALAN TAYLOR (1995). “Preference relations and measures in the context of fair division.” *Proceedings of the American Mathematical Society* 123, 2061-2070.
- BERNSTEIN, SERGEI (1917). “Attempt at an axiomatic foundation of probability theory [in Russian].” *Communications of the Kharkov Mathematical Society* [in Russian] 15, 209-274. Translation In: Sheynin, O. (2005). *Probability and Statistics: Russian Papers of the Soviet Period*. Berlin, Germany: NG Verlag.
- BRAMS, STEVEN AND ALAN TAYLOR (1995). “An envy-free cake division protocol.” *American Mathematical Monthly* 102, 9-18.
- BRAMS, STEVEN, ALAN TAYLOR, AND WILLIAM ZWICKER (1995). “Old and new moving-knife schemes.” *The Mathematical Intelligencer* 7, 30-35.
- CHATEAUNEUF, ALAIN (1985). “On the existence of a probability measure compatible with a total preorder on a Boolean algebra.” *Journal of Mathematical Economics* 14, 43-52.
- CRAMTON, PETER, ROBERT GIBBONS, AND PAUL KLEMPERER (1987). “Dissolving a partnership efficiently.” *Econometrica* 55, 615-632.
- CSEH, ÁGNES AND TAMÁS FLEINER (2020). “The complexity of cake cutting with unequal shares.” *ACM Transactions on Algorithms* 16, Article 29.
- DOWNEY, RODNEY AND MICHAEL FELLOWS (1999). *Parameterized Complexity*. New York, New York: Springer-Verlag.
- DUBINS, LESTER AND EDWIN SPANIER (1961). “How to cut a cake fairly.” *The American Mathematical Monthly* 68, 1-17.
- DE FINETTI, BRUNO (1937). “La prévision: ses lois logiques, ses sources subjectives [in French].” *Annales d’Institut Henri Poincaré* 7, 1-68.
- EDMONDS, JEFF AND KIRK PRUHS (2011). “Cake cutting really is not a piece of cake.” *ACM Transactions on Algorithms* 7, Article 51.
- ELIAZ, KFIR, RAN EILAT, AND XIAOSHENG MU (2021). “Bayesian privacy.” *Theoretical Economics* 16, 1557-1603.
- EVEN, SHIMON AND AZARIA PAZ (1984). “A note on cake-cutting.” *Discrete Applied Mathematics* 7, 285-296.
- FADEL, RONALD AND ILYA SEGAL (2009). “The communication cost of selfishness.” *Journal of Economic Theory* 144, 1895-1920.
- FAN, KY (1953). “Minimax theorems.” *Proceedings of the National Academy of Sciences of the United States of America* 39, 42-47.
- FLUM, JÖRG AND MARTIN GROHE (2006). *Parameterized Complexity Theory*. Berlin, Germany: Springer-Verlag.

- GAMOW, GEORGE AND MARVIN STERN (1958). *Puzzle-math*. New York, New York: Viking Press.
- GIBBARD, ALLAN (1973). “Manipulation of voting schemes: a general result.” *Econometrica* 41, 587-601.
- GRADWOHL, RONEN (2018). “Privacy in implementation.” *Social Choice and Welfare* 50, 547-580.
- GRADWOHL, RONEN AND RANN SMORODINSKY (2017). “Perception games and privacy.” *Games and Economic Behavior* 104, 293-308.
- HAKIMOV, RUSTAMDJAN AND MADHAV RAGHAVAN (2023). “Improving transparency and verifiability in school admissions: theory and experiment.” Working paper.
- HAUPT, ANDREAS AND ZOË HITZIG (2024). “Contextually private mechanisms.” Working paper.
- HAYEK, FRIEDRICH (1945). “The use of knowledge in society.” *The American Economic Review* 35, 519-530.
- KELLEY, JOHN L. (1959). “Measures on Boolean algebras.” *Pacific Journal of Mathematics* 9, 1165-1177.
- KOOPMAN, BERNARD (1940). “The axioms and algebra of intuitive probability.” *Annals of Mathematics* 41, 269-292.
- KRAFT, CHARLES, JOHN PRATT, AND ABRAHAM SEIDENBERG (1959). “Intuitive probability on finite sets.” *Annals of Mathematical Statistics* 30, 408-419.
- KUSHILEVITZ, EYAL AND NOAM NISAN (1997). *Communication Complexity*. Cambridge, United Kingdom: Cambridge University Press.
- LIAPOUNOFF, ALEKSANDR (1940). “Sur les fonctions-vecteurs complètement additives [in Russian].” *Izvestiya Akademii Nauk SSSR, Seriya Matematicheskaya* 4, 465-478.
- LOWRY, S. TODD (1987). *The Archaeology of Economic Ideas*. Durham, North Carolina: Duke University Press.
- MACKENZIE, ANDREW AND YU ZHOU (2022). “Menu mechanisms.” *Journal of Economic Theory* 204, Article 105511.
- MAGDON-ISMAIL, MALIK, COSTAS BUSCH, AND MUKKAI KRISHNAMOORTHY (2003). “Cake cutting is not a piece of cake.” In: *STACS 2003: 20th Annual Symposium on Theoretical Aspects of Computer Science*. Editors: Helmut Alt and Michel Habib. Berlin, Germany: Springer-Verlag.
- MARINACCI, MASSIMO (1993). “On the ranges of Baire and Borel measures, with applications to fine and tight comparative probabilities.” Manuscript.
- MAULDIN, RICHARD DANIEL (1981). *The Scottish Book: Mathematics from the Scottish Café*. Boston, Massachusetts: Birkhäuser.
- MCAVANEY, KEVIN, JACK ROBERTSON, AND WILLIAM WEBB (1992). “Ramsey partitions of integers and fair divisions.” *Combinatorica* 12, 193-201.
- MILGROM, PAUL AND ILYA SEGAL (2020). “Clock Auctions and Radio Spectrum Reallocation.” *Journal of Political Economy* 128, 1-31.
- MYERSON, ROGER AND MARK SATTERTHWAIT (1983). “Efficient mechanisms for bilateral trading.” *Journal of Economic Theory* 29, 265-281.
- NASH, JOHN (1951). “Non-cooperative games.” *Annals of Mathematics* 54, 286-295.
- VON NEUMANN, JOHN (1928). “Zur Theorie der Gesellschaftsspiele [in German].” *Mathematische Annalen* 100, 295-320. Translation In: Tucker, Albert and R. Duncan Luce (Editors) (1959). *Contributions to the Theory of Games, Volume IV*. Translator: Sonya Bargmann. Princeton, New Jersey: Princeton University Press.

- PROCACCIA, ARIEL (2009). “Thou shalt covet thy neighbor’s cake.” In: *IJCAI’09: Proceedings of the 21st International Joint Conference on Artificial Intelligence*. Editor: Hiroaki Kitano. San Francisco, California: Morgan Kaufmann Publishers Inc.
- PROCACCIA, ARIEL (2020). “Technical perspective: An answer to fair division’s most enigmatic question.” *Communications of the ACM* 63, 118.
- RICE, JOHN (1976). “The algorithm selection problem.” In: *Advances in Computers, Volume 15*. Editors: Morris Rubinoff and Marshall C. Yovits. New York, New York: Academic Press.
- ROBERTSON, JACK AND WILLIAM WEBB (1998). *Cake-Cutting Algorithms: Be Fair If You Can*. Boca Raton, Florida: CRC Press.
- RUDIN, WALTER (1976). *Principles of Mathematical Analysis (Third Edition)*. New York, NY: McGraw-Hill, Inc.
- SAVAGE, LEONARD (1954). *The Foundations of Statistics*. Wiley, NY: Dover Publications.
- SEGAL, ILYA (2007). “The communication requirements of social choice rules and supporting budget sets.” *Journal of Economic Theory* 136, 341-378.
- SEGAL, ILYA (2010). “Nash implementation with little communication.” *Theoretical Economics* 5, 51-71.
- SHISHIDO, HARUNORI AND DAO-ZHI ZENG (1999). “Mark-choose-cut algorithms for fair and strongly fair division.” *Group Decision and Negotiation* 8, 125-137.
- SION, MAURICE (1958). “On general minimax theorems.” *Pacific Journal of Mathematics* 8, 171-176.
- STEINHAUS, HUGO (1948). “The problem of fair division.” *Econometrica* 16, 101-104.
- WAKKER, PETER (1981). “Agreeing probability measures for comparative probability structures.” *The Annals of Statistics* 9, 658-662.
- WOEGINGER, GERHARD AND JIŘÍ SGALL (2007). “On the complexity of cake cutting.” *Discrete Optimization* 4, 213-220.
- WOODALL, DOUGLAS (1980). “Dividing a cake fairly.” *Journal of Mathematical Analysis and Applications* 78, 233-247.
- YAO, ANDREW (1979). “Some complexity questions related to distributive computing (Preliminary report).” In: *STOC ’79: Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing*. Atlanta, Georgia: Association for Computing Machinery.