The Computational Complexity of the Housing Market[*]

Edwin Lock[†]   Zephyr Qiu[‡]   Alexander Teytelboym[§]

February 13, 2024

Abstract

We prove that the classic problem of finding a competitive equilibrium in an exchange economy with indivisible goods, money, and unit-demand agents is PPAD-complete. In this "housing market", agents have preferences over the house and amount of money they end up with, but can experience income effects. Our results contrast with the existence of polynomial-time algorithms for related problems: Top Trading Cycles for the "housing exchange" problem in which there are no transfers and the Hungarian algorithm for the "housing assignment" problem in which agents' utilities are linear in money. Along the way, we prove that the Rainbow-KKM problem, a total search problem based on a generalization by Gale of the Knaster–Kuratowski–Mazurkiewicz lemma, is PPAD-complete. Our reductions also imply bounds on the query complexity of finding competitive equilibrium.

1 Introduction

The trade of indivisible objects (e.g., houses) among unit-demand agents is one of the classic settings in market design. If there is no divisible numéraire commodity such as money—we refer to such an economy as a "house exchange"—then Gale's Top Trading Cycles algorithm finds the unique allocation in the weak core Shapley and Scarf [1974], Roth and Postlewaite [1977]. If agents' utility functions are linear in money—we refer to such a transferable utility economy as a "housing assignment"—then an efficient allocation can be found by a straightforward solution (e.g., the Hungarian algorithm) to an assignment game [Koopmans and Beckmann, 1957, Shapley and Shubik, 1971]. However, both the assumptions of no transfers or of quasilinear utility are somewhat extreme, especially when one considers real-world trade of high-value objects such as houses. A more realistic assumption is that agents' utilities are non-linear in money, that is, agents experience income effects and their willingness to pay for a house might depend on their level of wealth (i.e., amount of money or the value of their own house). Surprisingly, a competitive equilibrium allocation always exists in such a "housing market", a result shown (under various assumptions) by Quinzii [1984], Gale [1984] and Svensson [1984]. But while the computational properties of the Top Trading Cycles algorithm in the house exchange model and the Hungarian algorithm in the house assignment model have been known for decades, nothing is known about the computational complexity of finding equilibrium in the housing market.

In this paper, we settle the computational complexity of the housing market in a general version introduced by Gale [1984]. In particular, we show that HOUSINGMARKET, the computational problem of finding an approximate competitive equilibrium in the housing market

model, is PPAD-complete [Papadimitriou, 1994] even when the market consists of three agents with identical preferences. Moreover, we show that the query complexity of HOUSINGMARKET with four or more agents is exponential in the approximation parameter. These hardness results stand in a stark contrast to the polynomial-time complexity of Top Trading Cycles algorithm for housing exchange and the Hungarian algorithm for housing assignment, and puts HOUSING-MARKET in same complexity class as the computation of Arrow-Debreu equilibria [Chen et al., 2009a] or Nash equilibria [Daskalakis et al., 2009, Chen et al., 2009b].[1]

We study the complexity of the housing market problem in two computational models. In the *white-box model*, also known as the polynomial function model, we assume that the preference functions of agents are given in the form of algorithms (or descriptions of Turing machines) with polynomial running time guarantees. We work in this model to show that HOUSINGMARKET is PPAD-complete. Secondly, we consider the *black-box model* (or *function oracle model*) in which the preference functions of agents are given by function oracles that can be queried to determine whether an agent demands a specific good at specific prices. The *query complexity* is then defined as the number of queries needed to find a solution. Note that the HOUSINGMARKET problem is at least as hard in the black-box model as in the white-box model, as the latter restricts the class of preferences to those expressible with polynomial-time functions. Moreover, the white-box model provides additional information about agent preferences through the algorithm implementing the preference functions. In the black-box model, we show exponential (in the approximation parameter) upper and lower bounds on the query complexity of HOUSINGMARKET with n agents.

The complexity class PPAD captures the computational complexity of many important problems that arise in economics and game theory. Two well-known PPAD-complete problems are SPERNER and CAKECUTTING. SPERNER is the problem of finding a panchromatic triangle in a triangulation of a simplex or hypercube, while CAKECUTTING captures the computational challenge of finding an envy-free allocation of a one-dimensional cake among players with heterogeneous preferences over pieces of cake. (See Sections 4.1, 4.2 and 5.2 for formal definitions of CAKECUTTING and SPERNER.) To prove our main result that HOUSINGMAR-KET is PPAD-complete, we first show that the problem is computationally equivalent to RAIN-BOWKKM, the generalization of the Knaster–Kuratowski–Mazurkiewicz lemma to multiple coverings due to Gale [1984]. Our polynomial-time reductions between the two problems also imply that HOUSINGMARKET and RAINBOWKKM have the same query complexity in the black-box model. We then develop polynomial-time reductions from two famous PPAD-complete problems (CAKECUTTING and 2D-SPERNER) to RAINBOWKKM to establish PPAD-hardness of RAINBOWKKM and HOUSINGMARKET. Our reduction from CAKECUTTING also establishes lower bounds on the query complexity of RAINBOWKKM. Surprisingly, our reduction from 2D-SPERNER implies that HOUSINGMARKET remains PPAD-hard even when the market consists of three agents and these agents have identical preferences. Finally, we show membership of RAINBOWKKM in PPAD by reducing RAINBOWKKM to SPERNER. This reduction uses the technique of labeling and coloring the vertices of a triangulated simplex initially developed in Su [1999] to provide the first computational procedure for finding an envy-free cake-cutting solution; this technique was then refined in Deng et al. [2012] using Kuhn's triangulation to prove PPAD-hardness of CAKECUTTING. Moreover, our reduction implies an exponential upper bound on the query complexity of RAINBOWKKM, which builds on a query complexity bound for SPERNER shown in Deng et al. [2011]. Our results are summarized in Table 1, and the main polynomial-time reductions developed in this paper are depicted in Fig. 1.

---

[1]See Part III in Rubinstein [2019] for a superb overview of PPAD and for many results about important economic problems, such as computing an approximate Bayes-Nash equilibrium and an approximate competitive equilibrium from equal incomes, in this class.
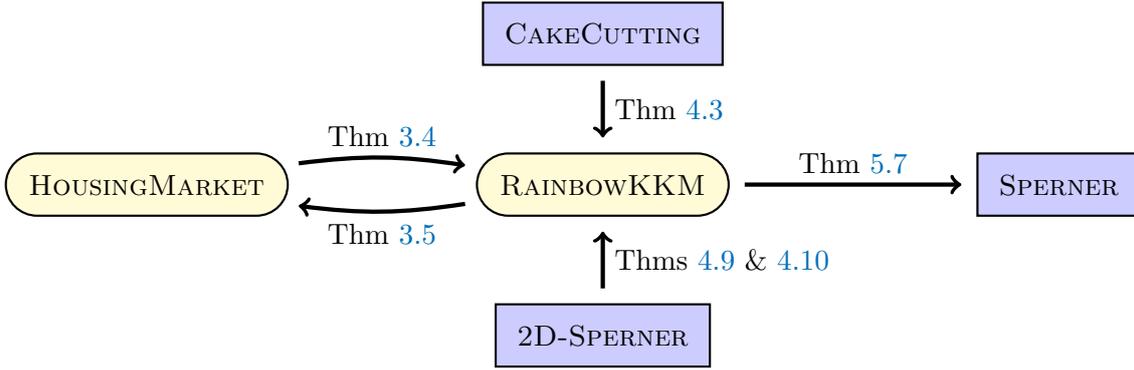
Figure 1: Illustration of the main reductions developed in this paper. The three problems shown in blue boxes are known to be PPAD-complete; CakeCutting was shown to be PPAD-complete by [Deng et al., 2012, Hollender and Rubinstein, 2023], 2D-Sperner in [Chen and Deng, 2009] and Sperner is a canonical PPAD-complete problem introduced in Papadimitriou [1994]. Arrows denote reductions from origin to destination.

Our main result has implications for economic theory, computation and market design more broadly, as the housing market model is a key building block in understanding competitive equilibrium with indivisible goods in the presence of income effects [Baldwin et al., 2023]. As the ability to compute equilibrium allocations is central to practical market design, our results suggest that more work is required to understand how to efficiently compute equilibrium allocation when budget constraints are a salient feature of the marketplace.

This paper relates to three strands of the literature. Firstly, our results add to a growing body of work showing the PPAD-completeness of computing equilibria in markets [Chen et al., 2009a, Chen and Teng, 2009]. Our model, which guarantees existence of equilibrium with indivisible goods and income effects, is conceptually related to the work of Chen et al. [2017] which shows PPAD-hardness of solving Arrow-Debreu markets with non-monotone utilities (over divisible goods). It is also similar in spirit (although not in techniques) to the model of Chen et al. [2022], which studies equilibria in the classic unit-demand pseudomarket of Hylland and Zeckhauser [1979]. However, we prove our results in Gale's model of the housing market which is very general and allows for externalities and non-monotonicity of utility in money. It is an open question whether the PPAD-completeness survives with stronger assumptions on preferences (as, e.g., in the Quinzii [1984] model).

Secondly, our reductions demonstrate parallels between CakeCutting and the HousingMarket and RainbowKKM problems. CakeCutting was first shown to be PPAD-complete by Deng et al. [2012] in a model with demand functions. PPAD-completeness of CakeCutting was then extended recently to the more stringent model with utility functions by Hollender and Rubinstein [2023]. We make use of this latter cake cutting model when we reduce CakeCutting to HousingMarket. In Gale's model of the housing market, agent preferences are also represented by demand functions. It is an open question whether PPAD-hardness continues to hold for HousingMarket when preferences are instead represented by utility functions.

Thirdly, our results are loosely related to a broader agenda on mechanism design with income effects (see, e.g., Che and Gale [1998], Benoit and Krishna [2001], Bhattacharya et al. [2010], Dobzinski et al. [2012], Saitoh and Serizawa [2008], Morimoto and Serizawa [2015], Baisa [2017]). These papers typically consider the design of optimal mechanisms in the presence of income effects, whereas we focus on computation of competitive equilibrium prices and allocations.

Complexity model	$n = 2$	$n = 3$	$n \geq 4$
white-box	P	PPAD-hard	PPAD-hard
black-box	$O(\log \frac{1}{\varepsilon})$	?	$\Theta(\log(\varepsilon))$

Table 1: The **computational complexity** (in the white-box model) and **query complexity** (in the black-box model) of the HOUSINGMARKET problem with n agents (and the computationally equivalent RAINBOWKKM problem). The two problems lie in PPAD for all n. The query complexity for $n = 3$ is unknown but expected to be $\Theta(\log(\varepsilon))$.

Organization. The rest of the paper is organised as follows. At the end of this section, we introduce the PPAD complexity class and notation. Section 2 states the model, defines competitive equilibrium, introduces the Rainbow-KKM lemma used by Gale [1984] to prove the existence of competitive equilibrium in his model. In this section, we also provide initial reductions. Section 3 establishes the computational equivalence of HOUSINGMARKET and RAINBOWKKM in the white-box and black-box models. In Section 4, we provide two proofs of the PPAD-hardness of HOUSINGMARKET and RAINBOWKKM with different implications, and derive a lower bound on the query complexity of the two problems. Section 5 proves membership of HOUSINGMARKET and RAINBOWKKM in PPAD, concluding our proof that the two problems are PPAD-complete, and gives an upper bound on the query complexity. Section 6 is a conclusion.

The PPAD complexity class. Computational problems for which the existence of a solution is guaranteed are called *total search problems*. The computational challenge is to find a solution. Examples of total search problems include integer factorisation, computing fixed points, and finding Nash equilibria. The class TFNP consists of all total search problems that can be solved in non-deterministic polynomial time; in other words, correctness of a solution to the problem can be verified in polynomial time. PPAD is the much-studied subclass of TFNP containing all problems that admit a polynomial-time reduction to the canonical problem ENDOFLINE [Papadimitriou, 1994]. For an introduction to PPAD and a definition of ENDOFLINE, we refer to [Goldberg, 2011]. A problem is PPAD-hard if ENDOFLINE reduces to it, and PPAD-complete if it lies in PPAD and is PPAD-hard. PPAD-hard problems are believed not to be solvable in polynomial time. As polynomial-time reductions are transitive, we can show membership of PPAD by reducing a problem to any other problem in PPAD, and PPAD-hardness by reducing from any problem that is PPAD-hard. Reductions between two problems also help us relate their query complexities.

Definition 1.1. For any two total search problems $\mathcal{P}$ and $\mathcal{Q}$, a *polynomial-time reduction* from $\mathcal{P}$ to $\mathcal{Q}$ consists of two polynomial-time computable functions f and g such that f maps any instance x of $\mathcal{P}$ to an instance $f(x)$ of $\mathcal{Q}$, and g maps any solution s for $f(\boldsymbol{x})$ to a solution $g(s)$ for $\boldsymbol{x}$.

Notation. $\mathbb{Z}$ denotes the integers, $\mathbb{N}$ the positive natural numbers, and $\mathbb{N}_0$ the natural numbers including 0. We write $[n] := \{1, \ldots, n\}$ and $[n]_0 := \{0, 1, \ldots, n\}$. Let $\Delta_{n-1} := \text{conv}\{\boldsymbol{e}^i \mid i \in [n]\}$ denote the standard $(n-1)$-simplex in $\mathbb{R}^n$. Throughout, the *distance between two points* $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$ is defined using the L_1 norm as $\|\boldsymbol{x} - \boldsymbol{y}\|_1 = \sum_{i \in [n]} |x_i - y_i|$. We write $\boldsymbol{x} \geq \boldsymbol{y}$ for two vectors $\boldsymbol{x}$ and $\boldsymbol{y}$ if the inequality holds elementwise. A function $f : X \to Y$ with $X, Y \subseteq \mathbb{R}^n$ is *Lipschitz-continuous with constant K* (or *K-Lipschitz*) if there exists a constant K such that $\|f(\boldsymbol{x}) - f(\boldsymbol{y})\|_1 \leq K\|\boldsymbol{x} - \boldsymbol{y}\|_1$ for all $\boldsymbol{x}, \boldsymbol{y} \in X$.

4

2 The Housing Market and Rainbow-KKM problems

We begin by introducing the HOUSINGMARKET and RAINBOWKKM problems and developing initial results.

2.1 The housing market

Consider a market with n agents $[n]$ that are each endowed with a house. Each house is identified with its agent, so the set of houses is $[n]$. The goal is to find a price for each house and an assignment of houses to agents so that every agent receives a house they prefer at these market prices.

The demand of each agent $i \in [n]$ is expressed using $n+1$ *preference sets* $P^i = (P^i_0, \ldots, P^i_n)$ that cover $\mathbb{R}^n$. At market prices $\boldsymbol{p} \in \mathbb{R}^n$, agent i demands good j if $\boldsymbol{p} \in P^i_j$, and nothing if $\boldsymbol{p} \in R^i_0$. If the agent demands nothing, she prefers to exchange her house for money without obtaining a new house. Note that an agent may be indifferent between multiple houses (or between demanding nothing and demanding one or more houses) if preference sets overlap at market prices. Gale [1984] imposes the following mild assumptions on the preference set of every agent $i \in [n]$. Let $B_n := \{\boldsymbol{p} \in \mathbb{R}^n \mid \boldsymbol{p} \geq 0, p_j = 0 \text{ for some } j \in [n]\}$ be the boundary of the positive quadrant. For each agent i, we assume that

(i) sets $P^i_0, P^i_1, \ldots, P^i_n$ are closed,

(ii) there exists some $M > 0$ such that $\boldsymbol{p} \notin P^i_j$ if $p_j \geq M$,

(iii) $P^i_1, \ldots, P^i_n$ cover $\Sigma^M_{n-1} = \{\boldsymbol{p} \in B_n \mid 0 \leq \boldsymbol{p} \leq M\}$.

Without loss of generality, we assume $M = 1$ (re-scaling prices if necessary) and write $\Sigma_n := \Sigma^M_n$.

Agents only wish to own at most one house, but otherwise the permissible preferences are very general: they allow for externalities among objects and demand to increase in prices (i.e., Giffen or Veblen goods). An *assignment* is a permutation of $[n]$ that maps agents to houses. A *market equilibrium* consists of market prices $\boldsymbol{p}$ and an assignment π such that every agent receives a house they demand at these prices; more formally, $\boldsymbol{p} \in P^i_{\pi(i)}$ for every agent $i \in [n]$. An ε-*equilibrium* is a pair $(\boldsymbol{p}, \pi)$ such that $\boldsymbol{p}$ is ε-close to $P^i_{\pi(i)}$ for every agent i, i.e. there exists $\boldsymbol{p}^i \in R^i_{\pi(i)}$ such that $\|\boldsymbol{p} - \boldsymbol{p}^i\|_1 \leq \varepsilon$ for every $i \in [n]$. Gale [1984] establishes that assumptions (i) to (iii) guarantee the existence of market equilibrium.

Theorem 2.1 ([Gale, 1984])**.** *Under assumptions (i) to (iii), there exists a market equilibrium.*

Hence finding an (approximate) equilibrium is a total search problem. As market prices can, in general, be irrational, we are interested in finding an ε-equilibrium for some approximation parameter ε. Moreover, we formalise preference sets in the computational setting by associating the preference sets $P^i = (P^i_1, \ldots, P^i_n)$ of each agent $i \in [n]$ with a *preference function* $f^i : \mathbb{R}^n \times [n]_0 \to \{0, 1\}$ so that $f^i(\boldsymbol{p}, j) := 1$ if $\boldsymbol{p} \in C^i_j$ and $f^i(\boldsymbol{p}, j) := 0$ otherwise. We say that a preference function satisfies assumptions (i) to (iii) if its associated preference sets do so. This allows us to define the computational problem as follows.

> **HousingMarket**
>
> **Input:** An approximation parameter $\varepsilon \in (0, \frac{1}{4})$. Preference functions $f^1, \ldots, f^n$ satisfying assumptions (i) to (iii) for the agents $[n]$.
> **Output:** An ε-equilibrium $(\boldsymbol{p}, \pi)$.

We refer to the computational problem with a fixed number n of agents as n-HousingMarket. In the black-box model, we assume that $f^1, \ldots, f^n$ are function oracles. In the white-box model, $f^1, \ldots, f^n$ are given as polynomial-time algorithms.

In Theorem 2.2, we show that n-HousingMarket reduces in polynomial time to $(n + 1)$-HousingMarket with the same approximation parameter. Starting with a market for n agents, the reduction adds an additional agent $n+1$ and corresponding house $n+1$. Preferences are designed so that, for any equilibrium prices $\boldsymbol{p} \in \mathbb{R}^{n+1}$ of the new market, agents $i \in [n]$ demand the same houses at $\boldsymbol{p}$ in the new market as they do at $(p_1, \ldots, p_n)$ in the original market. Hence the prices and allocation restricted to $[n]$ form an equilibrium for the original market. It is also clear from the reduction in the proof of Theorem 2.2 that the two problems have the same query complexity.

Theorem 2.2. *For any $n \geq 1$, there exists a polynomial-time reduction from n-HousingMarket to $(n + 1)$-HousingMarket.*

Proof. Suppose $(\varepsilon, f^1, \ldots, f^n)$ is an instance of n-HousingMarket associated with preference sets $P^1, \ldots, P^n$ for n agents. We construct an instance of $(n + 1)$-HousingMarket with the same approximation parameter and preference sets $Q^1, \ldots, Q^{n+1}$ for $n + 1$ agents as follows. For the first n agents $i \in [n]$, we extend the old preference sets P^i_j by defining

$$
\begin{aligned}
Q^i_0 &:= \mathbb{R}^n, \\
Q^i_j &:= \left\{ \boldsymbol{q} \in \mathbb{R}^{n+1} \mid (q_1, \ldots, q_n) \in P^i_j \right\} \text{ for all } j \in [n], \\
Q^i_{n+1} &:= \left\{ \boldsymbol{q} \in \mathbb{R}^{n+1} \mid q_{n+1} = 0 \right\}.
\end{aligned}
\tag{1}
$$

For the last agent $n + 1$, we define

$$
\begin{aligned}
Q^{n+1}_0 &:= \mathbb{R}^n, \\
Q^{n+1}_j &:= \left\{ \boldsymbol{q} \in \mathbb{R}^n \mid q_j = 0 \text{ and } q_{n+1} \geq \tfrac{3}{4} \right\} \text{ for all } j \in [n], \\
Q^{n+1}_{n+1} &:= \left\{ \boldsymbol{q} \in \mathbb{R}^n \mid q_{n+1} \leq \tfrac{3}{4} \right\}.
\end{aligned}
\tag{2}
$$

It is straightforward that the preference functions for $(Q^i_1, \ldots, Q^i_n)$ can be constructed efficiently and make at most one call to preference function f^i for $(P^i_1, \ldots, P^i_n)$. We argue that these preference sets satisfy Gale's assumptions (i) to (iii). It is immediate that Q^i_j is closed for all $i, j \in [n + 1]$, so (i) holds. Secondly, (ii) follows from the fact that P^i_j also satisfies (ii) when $i, j \in [n]$, and is immediate for $i = n + 1$ or $j = n + 1$. Finally, we show that (iii) holds. Indeed, suppose $\boldsymbol{q} \in \Sigma_{n+1}$. If $q_{n+1} = 0$, then $\boldsymbol{q} \in Q^i_{n+1}$ for all agents $i \in [n + 1]$. If $q_{n+1} > 0$, then $q_j = 0$ for some other house $j \in [n]$. For agents $i \in [n]$, this implies that $(q_1, \ldots, q_n) \in P^i_k$ for some $k \in [n]$ and so $\boldsymbol{q} \in Q^i_k$, as the preference sets $(P^i_1, \ldots, P^i_n)$ satisfy (iii). For agent $n + 1$, we see that $\boldsymbol{q} \in Q^{n+1}_j$ or $\boldsymbol{q} \in Q^{n+1}_{n+1}$.

Now let $(\boldsymbol{q}, \tau)$ be an ε-equilibrium of the new market with $n + 1$ agents. By definition, there exists $\boldsymbol{q}^i \in Q^i_{\tau(i)}$ with $\|\boldsymbol{q} - \boldsymbol{q}^i\|_1 \leq \varepsilon$ for every $i \in [n + 1]$. We argue that τ assigns house $n + 1$ to agent $n + 1$. Suppose not, so $\tau(n + 1) < n + 1$ and $\tau(i) = n + 1$ for some agent $i \in [n]$. By construction of Q^i_{n+1} we have $q^i_{n+1} = 0$. It follows that $q^{n+1}_{n+1} \leq 2\varepsilon \leq \tfrac{1}{2}$. But then $\boldsymbol{q}^{n+1} \notin Q^{n+1}_{\tau(n+1)}$, a contradiction. Define the permutation π of $[n]$ as the restriction of τ to $[n]$ and $\boldsymbol{p} := (q_1, \ldots, q_n)$. We now show that $(\boldsymbol{p}, \pi)$ is an ε-equilibrium for the original market with n agents. Let $\boldsymbol{p}^i := (q^i_1, \ldots, q^i_n)$ for each $i \in [n]$. We see that $\boldsymbol{q}^i \in Q^i_{\tau(i)}$ implies $\boldsymbol{p}^i \in P^i_{\pi(i)}$ by construction of Q^i_j. Moreover, $\|\boldsymbol{p} - \boldsymbol{p}^i\|_1 \leq \|\boldsymbol{q} - \boldsymbol{q}^i\| \leq \varepsilon$ for all $i \in [n]$. $\square$

2.2 The KKM and RainbowKKM problems

We now introduce the KKM and Rainbow-KKM lemmas, as well as their corresponding computational problems. The RainbowKKM problem will serve as our intermediary problem to show that HousingMarket is PPAD-complete. Recall the standard $(n-1)$-simplex $\Delta_{n-1} :=$ conv$\{e^i \mid i \in [n]\}$. For any $S \subseteq [n]$, we let $F_S := $ conv$\{e^i \mid i \in S\}$ be the face of Δ_{n-1} spanned by vectors e^i for all $i \in S$. A *KKM covering of* Δ_{n-1} is a collection of n closed subsets $C_1, \ldots, C_n$ of $\mathbb{R}^n$ such that $F_S \subseteq \bigcup_{i \in S} C_i$ for every $S \subseteq [n]$. The Knaster-Kuratowski-Mazurkiewicz (KKM) lemma states that every KKM covering $(C_1, \ldots, C_n)$ of Δ_{n-1} admits a point that is contained in all C_i, $i \in [n]$. For the proof of Theorem 2.1, Gale [1984] extends this result to families of KKM coverings.

Lemma 2.3 ([Knaster et al., 1929])**.** *Let* $(C_1, \ldots, C_n)$ *be a KKM covering of* Δ_{n-1}*. There exists a point* $\boldsymbol{x} \in \Delta_{n-1}$ *that lies in* $\bigcap_{i \in [n]} C_i$*.*

Lemma 2.4 ([Gale, 1984])**.** *Let* $C^1, \ldots, C^n$ *be* n *KKM coverings of* Δ_{n-1}*. There exists a permutation* π *of* $[n]$ *and a point* $\boldsymbol{x} \in \Delta_{n-1}$ *such that* $\boldsymbol{x} \in C^i_{\pi(i)}$ *for every* $i \in [n]$*.*

The proofs of Lemmas 2.3 and 2.4 do not lead to an efficient way to compute the point $\boldsymbol{x}$ in question. Moreover, as $\boldsymbol{x}$ can be irrational, we consider the computational problems of finding an approximate solution. To formalize KKM coverings in our computational models, we associate a KKM covering $(C_1, \ldots, C_n)$ of Δ_{n-1} with a *KKM covering function* $g : \mathbb{R}^n \times [n] \to \{0, 1\}$ by defining $g(\boldsymbol{x}, j) := 1$ if $\boldsymbol{x} \in C^i_j$ and $g^i(\boldsymbol{x}, j) := 0$ otherwise. This leads to the following computational problems in which the covering functions are given as function oracles or as polynomial-time algorithms.

KKM

Input: Approximation parameter $\varepsilon \in (0, \frac{1}{4})$. Dimension n. KKM covering function g of Δ_{n-1}.

Output: Point $\boldsymbol{x}$ for which $\|\boldsymbol{x} - \boldsymbol{x}^i\|_1 \leq \varepsilon$ for some $\boldsymbol{x}^i \in C_i$ for every $i \in [n]$.

We say that a pair $(\boldsymbol{x}, \pi)$ consisting of point $\boldsymbol{x} \in \Delta_{n-1}$ and permutation π of $[n]$ is an ε*-approximate Rainbow-KKM solution* if there exists $\boldsymbol{x}^i \in C^i_{\pi(i)}$ with $\|\boldsymbol{x} - \boldsymbol{x}^i\|_1 \leq \varepsilon$ for every $i \in [n]$. We refer to the KKM problem with fixed dimension $n = 3$ as 3D-KKM.

RainbowKKM

Input: Approximation parameter $\varepsilon \in (0, \frac{1}{4})$. KKM covering functions $g^1, \ldots, g^n$ of Δ_{n-1}.

Output: ε-approximate Rainbow-KKM solution $(\boldsymbol{x}, \pi)$.

We refer to the RainbowKKM problem with fixed dimension n as n-RainbowKKM. As we will see in Section 4.2, finding an approximate Rainbow-KKM point is at least as hard as finding an approximate KKM point.

2.3 Sparse KKM coverings

In Section 5, we reduce RainbowKKM to the problem Sperner. In the reduction, we assume that each KKM covering $(C_1, \ldots, C_n)$ of Δ_{n-1} satisfies $C_i \cap F_{[n] \setminus \{i\}} = \emptyset$ for every $i \in [n]$. We call such a KKM covering *sparse*. We now argue that this is without loss of generality, as we can reduce any RainbowKKM instance to an instance with sparse KKM coverings. Note that

we cannot simply make a KKM covering $(C_1, \ldots, C_n)$ sparse by subtracting $F_{[n]\setminus\{i\}}$ from C_i for each $i \in [n]$, as the resulting sets are not closed.

Suppose $(C_1, \ldots, C_n)$ is a KKM covering of Δ_{n-1}. We now construct a KKM covering $(D_1, \ldots, D_n)$ parametrised by $\delta > 0$ that is sparse. Moreover, we will see that for every $\boldsymbol{x} \in D_i$ there exists a point $\boldsymbol{y} \in C_i$ close to $\boldsymbol{x}$. This forms the basis of our reduction from RAINBOWKKM to the version with sparse KKM coverings.

Fix parameter $\delta \in (0, \frac{1}{4})$. First we define a projection function $\tau : \Delta_{n-1} \to \Delta_{n-1}$. Let $\tau(\boldsymbol{x})$ be the vector obtained from $\boldsymbol{x}$ by reducing all entries less or equal to δ down to 0, and normalizing the remaining entries so that the entries sum to 1. We write

$$\hat{\tau}(\boldsymbol{x})_i := \begin{cases} 0 & \text{if } x_i \leq \delta, \\ x_i & \text{else,} \end{cases}$$

and $\tau(\boldsymbol{x}) = \frac{1}{\|\hat{\tau}(\boldsymbol{x})\|_1} \hat{\tau}(\boldsymbol{x})$. Note that $\|\tau(\boldsymbol{x}) - \boldsymbol{x}\|_1 \leq 2n\delta$ for any $\boldsymbol{x} \in \Delta_{n-1}$. We define D_i for each $i \in [n]$ as

$$D_i := \{\boldsymbol{x} \in \Delta_{n-1} \mid x_i \geq \delta \text{ and } \tau(\boldsymbol{x}) \in C_i\}. \tag{3}$$

We now argue that $(D_1, \ldots, D_n)$ is a sparse KKM covering of Δ_{n-1}. To see that each D_i is closed, write D_i as the union of sets $D_i^T := \{\boldsymbol{x} \in D_i \mid \{j \in [n] \mid x_j \leq \delta\} = T\}$ for all $T \subseteq [n] \setminus \{i\}$. By a standard limiting argument and the fact that C_i is closed, we see that each D_i^T is closed, and so their union D_i is also closed. Next, we see that $(D_1, \ldots, D_n)$ satisfies the KKM covering property. Fix $S \subseteq [n]$ and $\boldsymbol{x} \in F_S$. We let $S' = \{i \in [n] \mid \tau(\boldsymbol{x})_i > 0\}$, so $S' \subseteq S$ and $\tau(\boldsymbol{x}) \in F_{S'}$. As $(C_1, \ldots, C_n)$ is a KKM covering, $\tau(\boldsymbol{x}) \in \bigcup_{i \in S'} C_i$. Fix $j \in S'$ with $\tau(\boldsymbol{x}) \in C_j$. Note that $\tau(\boldsymbol{x})_j > 0$ implies $\tau(\boldsymbol{x})_j > \delta$. So, by construction, $\tau(\boldsymbol{x}) \in D_j$. Hence $\tau(\boldsymbol{x}) \in \bigcup_{i \in S} D_i$, which concludes the proof that $(D_1, \ldots, D_n)$ is a KKM covering. Finally, as D_i contains only points $\boldsymbol{x}$ satisfying $x_i \geq \delta$ and $\boldsymbol{x} \in D_{[n]\setminus\{i\}}$ implies $x_i = 0$, we see that $(D_1, \ldots, D_n)$ is sparse.

Proposition 2.5. *There exists a polynomial-time reduction from n-RAINBOWKKM to the problem $n - \text{RAINBOWKKM}$ with sparse coverings.*

Proof. Suppose $(\varepsilon, g^1, \ldots, g^n)$ is an instance of n-RAINBOWKKM associated with KKM coverings $C^1, \ldots, C^n$. We construct an instance $(\varepsilon' = \frac{\varepsilon}{2}, h^1, \ldots, h^n)$ of n-RAINBOWKKM associated with sparse KKM coverings $D^1, \ldots, D^n$ that are constructed as described in (3) with parameter $\delta = \frac{\varepsilon}{8n}$. It is straightforward that h^i can be implemented efficiently with at most one call to g^i.

Suppose $(\boldsymbol{x}, \pi)$ is a solution for instance $(\varepsilon', h^1, \ldots, h^n)$. Hence there exists $\boldsymbol{x}^i \in D_{\pi(i)}^i$ with $\|\boldsymbol{x} - \boldsymbol{x}^i\|_1 \leq \varepsilon'$ for every $i \in [n]$. By construction of the sets D_j^i in (3), we have $\tau(\boldsymbol{x}^i) \in C_{\pi(i)}^i$. Moreover, we know that $\|\boldsymbol{x} - \tau(\boldsymbol{x})\|_1 \leq 2n\delta$ and $\|\boldsymbol{x}^i - \tau(\boldsymbol{x}^i)\|_1 \leq 2n\delta$, so $\|\tau(\boldsymbol{x}) - \tau(\boldsymbol{x}^i)\|_1 \leq \varepsilon' + 4n\delta = \varepsilon$. Hence $(\tau(\boldsymbol{x}), \pi)$ is a solution for the original instance $(\varepsilon, g^1, \ldots, g^n)$. $\square$

3 Reductions between HousingMarket and RainbowKKM

We now establish the computational equivalence of HOUSINGMARKET and RAINBOWKKM by proving polynomial-time reductions in both directions. We note that our reductions preserve the dimensionality of the problems; that is, an instance of n-HOUSINGMARKET is reduced to an instance of n-RAINBOWKKM, and vice versa. Moreover, it follows from our reductions that n-HOUSINGMARKET and n-RAINBOWKKM have the same query complexity.

In Section 3.1, we define a homeomorphism $\phi : \Sigma_n \to \Delta_{n-1}$ from the 'relevant' domain Σ_n of the housing market to the domain Δ_{n-1} of the Rainbow-KKM problem. This homeomorphism was first introduced by Gale [1984] to prove equilibrium existence for the housing market via the Rainbow-KKM lemma. Section 3.2 then provides the reductions between HOUSINGMARKET and RAINBOWKKM.

8

3.1 Connecting the domains of HousingMarket and RainbowKKM

Recall that $\Sigma_n = \{\boldsymbol{x} \in \mathbb{R}^n \mid 0 \le \boldsymbol{x} \le 1, \text{ and } x_i = 0 \text{ for some } i \in [n]\}$ is the intersection of the unit cube with the boundary of the positive orthant B_n. For each permutation π of $[n]$, define the two sets

$$\Sigma_\pi := \{\boldsymbol{p} \in \Sigma_n \mid p_{\pi(1)} \ge p_{\pi(2)} \ge \cdots \ge p_{\pi(n)} = 0\}$$
$$\text{and } \Delta_\pi := \{\boldsymbol{x} \in \Delta_{n-1} \mid x_{\pi(1)} \le x_{\pi(2)} \le \cdots \le x_{\pi(n)}\}.$$

It is immediate that $\Sigma_n = \bigcup_\pi \Sigma_\pi$ and $\Delta_{n-1} = \bigcup_\pi \Delta_\pi$. For each $\boldsymbol{p} \in \Sigma_\pi$, we define

$$\phi(\boldsymbol{p})_{\pi(k)} := \frac{1 - p_{\pi(1)}}{n} + \frac{p_{\pi(1)} - p_{\pi(2)}}{n - 1} + \frac{p_{\pi(k-1)} - p_{\pi(k)}}{n - k + 1} \tag{4}$$

The inverse $\phi^{-1} : \Delta_{n-1} \to \Sigma_n$ of ϕ is defined, for every $\boldsymbol{x} \in \Delta_\pi$, as

$$\phi^{-1}(\boldsymbol{x})_{\pi(k)} := 1 - x_{\pi(1)} - x_{\pi(2)} - \cdots - x_{\pi(k-1)} - (n - k + 1)x_{\pi(k)}. \tag{5}$$

Lemma 3.1 shows that ϕ^{-1} is indeed the inverse of ϕ. Moreover, Lemmas 3.2 and 3.3 prove that ϕ is n-Lipschitz and ϕ^{-1} is n^2-Lipschitz. In particular, this establishes that ϕ is a homeomorphism.

Lemma 3.1. *The function ϕ is a bijection from Σ_n to Δ_{n-1}, and ϕ^{-1} is its inverse. Moreover, $\phi(\Delta_\pi) = \Sigma_\pi$ for every permutation π of $[n]$.*

Proof. It is straightforward to check that $\phi^{-1}(\phi(\boldsymbol{p})) = \boldsymbol{p}$ for every $\boldsymbol{p} \in \Sigma_n$ and $\phi(\phi^{-1}(\boldsymbol{x})) = \boldsymbol{x}$ for every $\boldsymbol{x} \in \Delta_{n-1}$, so ϕ^{-1} is the inverse of ϕ.

Next we show that $\phi(\Sigma_\pi) = \Delta_\pi$. Fix $\boldsymbol{p} \in \Sigma_\pi$. By construction of Σ_π, we have $p_{\pi(k-1)} - p_{\pi(k)} \ge 0$ for all $k \in \{2, \ldots, n\}$. It follows from (4) that $\phi(\boldsymbol{p})_{\pi(1)} \ge 0$ and $\phi(\boldsymbol{p})_{\pi(k-1)} \le \phi(\boldsymbol{p})_{\pi(k)}$ for all $k \in \{2, \ldots, n\}$, and

$$\sum_{i \in [n]} \phi(\boldsymbol{p})_{\pi(i)} = n \cdot \frac{1 - p_{\pi(1)}}{n} + (n-1) \cdot \frac{p_{\pi(1)} - p_{\pi(2)}}{n-1} + \cdots + 1 \cdot \frac{p_{\pi(n-1)} - p_{\pi(n)}}{1} = 1 - p_{\pi(n)} = 1.$$

Hence, $\phi(\boldsymbol{p}) \in \Delta_\pi$. As $\boldsymbol{p}$ was chosen arbitrarily, we get $\phi(\Sigma_\pi) \subseteq \Delta_\pi$. To show that $\phi(\Sigma_\pi) = \Delta_\pi$, we now fix $\boldsymbol{x} \in \Delta_\pi$ and show that $\phi^{-1}(\boldsymbol{x}) \in \Sigma_\pi$. Let $\boldsymbol{p} := \phi^{-1}(\boldsymbol{x})$. By (5), we have $p_{\pi(k-1)} - p_{\pi(k)} = (n-k+1)(x_{\pi(k-1)} - x_{\pi(k)}) \le 0$, and so $p_{\pi(k-1)} \ge p_{\pi(k)}$, for every $k \in \{2, \ldots, n\}$. Moreover, $\boldsymbol{x} \in \Delta_\pi$ tells us that $x_1 \le \frac{1}{n}$ and $\sum_{i \in [n]} x_{\pi(i)} = 1$. Hence, $p_{\pi(1)} = 1 - nx_{\pi(1)} \in [0, 1]$ and $p_{\pi(n)} = 1 - \sum_{i \in [n]} x_{\pi(i)} = 0$. $\square$

Lemma 3.2. *The function ϕ is n-Lipschitz.*

Proof. In order to prove that ϕ is n-Lipschitz, we need to show $\|\phi(\boldsymbol{p}) - \phi(\boldsymbol{q})\|_1 \le n\|\boldsymbol{p} - \boldsymbol{q}\|_1$ for all $\boldsymbol{p}, \boldsymbol{q} \in \Sigma_n$. Fix two points $\boldsymbol{p}, \boldsymbol{q} \in \Sigma_n$, and let $\|\boldsymbol{p} - \boldsymbol{q}\|_1 = \delta$. If $\boldsymbol{p}$ and $\boldsymbol{q}$ both lie in Σ_π for some permutation π, then

$$|\phi(\boldsymbol{p})_{\pi(k)} - \phi(\boldsymbol{q})_{\pi(k)}| = \sum_{j \in [k]} \left(\frac{1}{n - j + 1} - \frac{1}{n - j + 2} \right) |p_{\pi(j)} - q_{\pi(j)}| \le \sum_{j \in [k]} |p_{\pi(j)} - q_{\pi(j)}| \le \delta \tag{6}$$

and so $\|\phi(\boldsymbol{p}) - \phi(\boldsymbol{q})\|_1 \le n\delta$.

From now on, we assume that $\boldsymbol{p} \in \Sigma_\pi$ and $\boldsymbol{q} \notin \Sigma_\pi$. Suppose first that $\boldsymbol{p}$ and $\boldsymbol{q}$ differ only in one entry, so $\boldsymbol{q} = \boldsymbol{p} + \alpha \boldsymbol{e}^{\pi(j)}$ for some $j \in [n]$. Without loss of generality, we can assume that $\alpha > 0$ (by relabeling $\boldsymbol{p}$ and $\boldsymbol{q}$ if $\alpha < 0$). Note that $\boldsymbol{q} \notin \Sigma_\pi$ implies $j \ge 2$. Moreover, increasing

9

$p_{\pi(j)}$ by α moves it to the position of $p_{\pi(i)}$ for some $i < j$, and moves entries $p_{\pi(i)}, \ldots, p_{\pi(j-1)}$ one position to the right. The entries of $\boldsymbol{p}$ and $\boldsymbol{q}$ in descending order are, thus, given by

$$p_{\pi(1)}, p_{\pi(2)}, \ldots, p_{\pi(i-1)}, p_{\pi(i)}, \ldots, p_{\pi(j-1)}, p_{\pi(j)}, p_{\pi(j+1)}, \ldots, p_{\pi(n)}, \tag{7}$$

$$p_{\pi(1)}, p_{\pi(2)}, \ldots, p_{\pi(i-1)}, p_{\pi(j)} + \alpha, p_{\pi(i)}, \ldots, p_{\pi(j-1)}, p_{\pi(j+1)}, \ldots, p_{\pi(n)}. \tag{8}$$

Observe that

$$p_{\pi(j)} \le p_{\pi(k)} \le p_{\pi(j)} + \alpha, \ \forall k \in \{i, \ldots, j\}. \tag{9}$$

We will now show that $|\phi(\boldsymbol{p})_{\pi(k)} - \phi(\boldsymbol{q})_{\pi(k)}| \le \alpha$ for all $k \in [n]$. Fix some $k \in [n]$ and note

$$\phi(\boldsymbol{p})_{\pi(k)} = \frac{1 - p_{\pi(1)}}{n} + \sum_{l=2}^{k} \frac{p_{\pi(l-1)} - p_{\pi(l)}}{n - l + 1}.$$

If $k \in \{1, \ldots, i-1\}$, then $\phi(\boldsymbol{p})_{\pi(k)} = \phi(\boldsymbol{q})_{\pi(k)}$ by (7). Now suppose that $k \in \{i, \ldots, j-1\}$. Then

$$\phi(\boldsymbol{q})_{\pi(k)} = \frac{1 - p_{\pi(1)}}{n} + \sum_{l=2}^{i-1} \frac{p_{\pi(l-1)} - p_{\pi(l)}}{n - l + 1}$$

$$+ \frac{p_{\pi(i-1)} - (p_{\pi(j)} + \alpha)}{n - i + 1} + \frac{(p_{\pi(j)} + \alpha) - p_{\pi(i)}}{n - i + 2}$$

$$+ \sum_{l=i+1}^{k} \frac{p_{\pi(l-1)} - p_{\pi(l)}}{n - l + 2},$$

and so

$$|\phi(\boldsymbol{p})_{\pi(k)} - \phi(\boldsymbol{q})_{\pi(k)}| = \frac{p_{\pi(i-1)} - p_{\pi(i)}}{n - i + 1} - \frac{p_{\pi(i-1)} - (p_{\pi(j)} + \alpha)}{n - i + 1} - \frac{p_{\pi(j)} + \alpha - p_{\pi(i)}}{n - i + 2}$$

$$+ \sum_{l=i+1}^{k} \left(\frac{1}{n - i + 1} - \frac{1}{n - i + 2} \right) (p_{\pi(l-1)} - p_{\pi(l)})$$

$$\le \left(\frac{1}{n - i + 1} - \frac{1}{n - i + 2} \right) (p_{\pi(j)} + \alpha - p_{\pi(i)} + p_{\pi(i)} - p_{\pi(k)})$$

$$\le p_{\pi(j)} + \alpha - p_{\pi(k)} \le \alpha.$$

The last inequality follows from (9). Next, suppose $k = j$. In this case, we have

$$\phi(\boldsymbol{q})_{\pi(j)} = \frac{1 - p_{\pi(1)}}{n} + \sum_{l=2}^{i-1} \frac{p_{\pi(l-1)} - p_{\pi(l)}}{n - l + 1} + \frac{p_{\pi(i-1)} - (p_{\pi(j)} + \alpha)}{n - i + 1},$$

so

$$|\phi(\boldsymbol{p})_{\pi(j)} - \pi(\boldsymbol{q})_{\pi(j)}| = \sum_{l=i}^{j} \frac{p_{\pi(l-1)} - p_{\pi(l)}}{n - l + 1} - \frac{p_{\pi(i-1)} - (p_{\pi(j)} + \alpha)}{n - i + 1}$$

$$= \frac{p_{\pi(j)} + \alpha - p_{\pi(i)}}{n - i + 1} + \sum_{l=i+1}^{j} \frac{p_{\pi(l-1)} - p_{\pi(l)}}{n - l + 1}$$

$$\le p_{\pi(j)} + \alpha - p_{\pi(i)} + p_{\pi(i)} - p_{\pi(j)} = \alpha.$$

10

Now suppose that $k \in \{j+1, \ldots, n\}$. Hence,

$$\phi(\boldsymbol{q})_{\pi(k)} = \frac{1 - p_{\pi(1)}}{n} + \sum_{l=2}^{i-1} \frac{p_{\pi(l-1)} - p_{\pi(l)}}{n - l + 1}$$

$$+ \frac{p_{\pi(i-1)} - (p_{\pi(j)} + \alpha)}{n - i + 1} + \frac{p_{\pi(j)} + \alpha - p_{\pi(i)}}{n - i + 2}$$

$$+ \sum_{l=i}^{j-1} \frac{p_{\pi(l-1)} - p_{\pi(l)}}{n - l + 1} + \frac{p_{\pi(j-1)} - p_{\pi(j+1)}}{n - j + 2}$$

$$+ \sum_{l=j+2}^{k} \frac{p_{\pi(l-1)} - p_{\pi(l)}}{n - l + 1}.$$

So we have

$$|\phi(\boldsymbol{p})_{\pi(k)} - \phi(\boldsymbol{q})_{\pi(k)}| = \frac{p_{\pi(i-1)} - p_{\pi(i)}}{n - i + 1} - \frac{p_{\pi(i-1)} - p_{\pi(j)} + \alpha}{n - i + 1} - \frac{p_{\pi(j)} + \alpha - p_{\pi(i)}}{n - i + 2}$$

$$+ \sum_{l=i+1}^{j-1} \left(\frac{1}{n - j + 1} - \frac{1}{n - j + 2} \right) (p_{\pi(l-1)} - p_{\pi(l)})$$

$$+ \frac{p_{\pi(j-1)} - p_{\pi(j)}}{n - j + 1} - \frac{p_{\pi(j)} - p_{\pi(j+1)}}{n - i + 2} - \frac{p_{\pi(j-1)} - p_{\pi(j+1)}}{n - i + 2}$$

$$= \left(\frac{1}{n - j + 1} - \frac{1}{n - j + 2} \right) (p_{\pi(j)} + \alpha - p_{\pi(i)})$$

$$+ \left(\frac{1}{n - j + 1} - \frac{1}{n - j + 2} \right) (p_{\pi(i)} - p_{\pi(j-1)})$$

$$+ \left(\frac{1}{n - j + 1} - \frac{1}{n - j + 2} \right) (p_{\pi(j-1)} - p_{\pi(j)})$$

$$\leq \alpha.$$

Finally, suppose $\boldsymbol{p}$ and $\boldsymbol{q}$ differ in arbitrarily many entries, so $\boldsymbol{q} = \boldsymbol{p} + \sum_{i \in [n]} \alpha_i \boldsymbol{e}^i$ for some $\alpha \in \mathbb{R}^n$ and we have $\delta = \|\boldsymbol{\alpha}\|_1$. Define $\boldsymbol{q}^{(k)} = \boldsymbol{p} + \sum_{i=1}^{k} \alpha_i \boldsymbol{e}^i$. (So $\boldsymbol{q}^{(0)} = \boldsymbol{p}$ and $\boldsymbol{q}^{(n)} = \boldsymbol{q}$.) Then

$$\|\phi(\boldsymbol{p}) - \phi(\boldsymbol{q})\|_1 = \left\| \sum_{i=1}^{n} \phi(\boldsymbol{q}^{(i-1)}) - \phi(\boldsymbol{q}^{(i)}) \right\|_1 \leq \sum_{i=1}^{n} \left\| \phi(\boldsymbol{q}^{(i-1)}) - \phi(\boldsymbol{q}^{(i)}) \right\|_1 \leq n \sum_{i=1}^{n} |\alpha_i| \leq n\delta.$$

The first inequality makes use of the triangle inequality. The second inequality applies the result from the previous paragraph to each $\|\phi(\boldsymbol{p}^{(i-1)}) - \phi(\boldsymbol{q}^{(i)})\|_1$, as $\boldsymbol{q}^{(i-1)}$ and $\boldsymbol{q}^{(i)}$ differ in one entry, by α_i. $\qquad\square$

Lemma 3.3. *The function ϕ^{-1} is n^2-Lipschitz.*

Proof. Define $\mathbb{R}^n_\pi = \{\boldsymbol{x} \in \mathbb{R}^n \mid x_{\pi(1)} \leq x_{\pi(2)} \leq \cdots \leq x_{\pi(n)}\}$ for each permutation π of $[n]$. For the purpose of this proof, we extend ϕ^{-1} to domain $\mathbb{R}^n$ by defining $\phi^{-1}(\boldsymbol{x})_{\pi(k)}$ as in (5) for every $\boldsymbol{x} \in \mathbb{R}^n_\pi$. We will show that $\|\phi^{-1}(\boldsymbol{x}) - \phi^{-1}(\boldsymbol{y})\| \leq n^2 \|\boldsymbol{x} - \boldsymbol{y}\|_1$ for all $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$. The proof is conceptually similar to that of Lemma 3.2. Fix $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$, and define $\delta := \|\boldsymbol{x} - \boldsymbol{y}\|_1$. If $\boldsymbol{x}$ and $\boldsymbol{y}$ both lie in $\mathbb{R}^n_\pi$ for some permutation π, then, by applying the definition (5) of ϕ^{-1}, we have

$$|\phi^{-1}(\boldsymbol{x})_{\pi(k)} - \phi^{-1}(\boldsymbol{y})_{\pi(k)}| \leq (n-k+1)|x_{\pi(k)} - y_{\pi(k)}| + \sum_{i=1}^{k-1} |x_{\pi(i)} - y_{\pi(i)}| \leq (n-k+1)\delta + (k-1)\delta = n\delta,$$

11

so $\|\phi^{-1}(\boldsymbol{x}) - \phi^{-1}(\boldsymbol{y})\|_1 \leq n^2\delta$.

From now on, we assume that $\boldsymbol{x} \in \mathbb{R}^n_\pi$ and $\boldsymbol{y} \notin \mathbb{R}^n_\pi$. Suppose first that $\boldsymbol{p}$ and $\boldsymbol{q}$ differ only in one entry, so $\boldsymbol{q} = \boldsymbol{p} - \alpha e^{\pi(j)}$ for some $j \in [n]$. Without loss of generality, assume $\alpha > 0$ (relabeling $\boldsymbol{p}$ and $\boldsymbol{q}$ if necessary). Note that $\boldsymbol{q} \notin \mathbb{R}^n_\pi$ implies $j \geq 2$. Decreasing entry $x_{\pi(j)}$ by α moves it to the position of $x_{\pi(i)}$ for some $i < j$ (strict inequality holds due to $\boldsymbol{y} \notin \mathbb{R}^n_\pi$), and moves entries $x_{\pi(i)}, \ldots, x_{\pi(j-1)}$ one position to the right. The entries of $\boldsymbol{y}$ in ascending order are, thus,

$$x_{\pi(1)}, x_{\pi(2)}, \ldots, x_{\pi(i-1)}, x_{\pi(j)} - \alpha, x_{\pi(i)}, \ldots, x_{\pi(j-1)}, x_{\pi(j+1)}, \ldots, x_{\pi(n)}. \tag{10}$$

Observe that

$$x_{\pi(j)} - \alpha \leq x_{\pi(k)} \leq x_{\pi(j)}, \ \forall k \in \{i, \ldots, j\}. \tag{11}$$

We will now show that $|\phi^{-1}(\boldsymbol{x})_{\pi(k)} - \phi^{-1}(\boldsymbol{y})_{\pi(k)}| \leq n\alpha$ for all $k \in [n]$, which implies $\|\phi^{-1}(\boldsymbol{x}) - \phi^{-1}(\boldsymbol{y})\|_1 \leq n^2\delta$. Fix $k \in [n]$ and note that

$$\phi^{-1}(\boldsymbol{x})_{\pi(k)} = 1 - \sum_{l=1}^{k-1} x_{\pi(l)} - (n - k + 1)x_{\pi(k)}.$$

We distinguish between four cases. If $k \in \{1, \ldots, i-1\}$, then $\boldsymbol{x}$ and $\boldsymbol{y}$ coincide on entries $\pi(1), \ldots, \pi(k)$, so $\phi^{-1}(\boldsymbol{x})_{\pi(k)} = \phi(\boldsymbol{y})_{\pi(k)}$. Secondly, suppose $k \in \{i, \ldots, j-1\}$. Then $\phi^{-1}(\boldsymbol{y})_{\pi(k)} = 1 - \sum_{l=1}^{k-1} x_{\pi(l)} - (x_{\pi(j)} - \alpha) - (n-k)\boldsymbol{x}_{\pi(k)}$, which implies $|\phi^{-1}(\boldsymbol{x})_{\pi(k)} - \phi^{-1}(\boldsymbol{y})_{\pi(k)}| \leq |x_{\pi(j)} - \alpha - x_{\pi(k)}| \leq \alpha$ by (11). Thirdly, suppose $k = j$. Then $\phi^{-1}(\boldsymbol{y})_{\pi(k)} = 1 - \sum_{l=1}^{i-1} x_{\pi(l)} - (n - i + 1)(x_{\pi(j)} - \alpha)$, so $|\phi^{-1}(\boldsymbol{x})_{\pi(k)} - \phi^{-1}(\boldsymbol{y})_{\pi(k)}| \leq \left| \sum_{l=i}^{j-1} (x_{\pi(j)} - x_{\pi(l)}) - (n - i + 1)\alpha \right| \leq n\alpha$, by (11). Lastly, suppose $k \in \{j+1, \ldots, n\}$. Then $\phi^{-1}(\boldsymbol{y})_{\pi(k)} = 1 + \alpha - \sum_{l=1}^{k-1} x_{\pi(l)} - (n-k+1)x_{\pi(k)}$, so $|\phi^{-1}(\boldsymbol{x})_{\pi(k)} - \phi^{-1}(\boldsymbol{y})_{\pi(k)}| = \alpha$.

Now suppose $\boldsymbol{x}$ and $\boldsymbol{y}$ differ in arbitrarily many entries. As in the previous proof, we have $\boldsymbol{y} = \boldsymbol{x} + \sum_{i=1}^n \alpha_i e^i$ for some $\boldsymbol{\alpha} \in \mathbb{R}^n$ and $\delta = \|\boldsymbol{\alpha}\|_1$. Define $\boldsymbol{x}^{(k)} = \boldsymbol{x} + \sum_{i=1}^k \alpha_i e^i$ for all $k \in \{0, \ldots, n\}$, so $\boldsymbol{x}^{(0)} = \boldsymbol{x}$ and $\boldsymbol{x}^{(n)} = \boldsymbol{y}$. Then $\boldsymbol{x}^{(i-1)}$ and $\boldsymbol{x}^{(i)}$ differ on one entry, by α_i. Hence,

$$\|\phi^{-1}(\boldsymbol{x}) - \phi^{-1}(\boldsymbol{y})\|_1 \leq \sum_{i=1}^n \|\phi^{-1}(\boldsymbol{x}^{(i-1)}) - \phi^{-1}(\boldsymbol{x}^{(i)})\|_1 \leq n^2 \sum_{i=1}^n |\alpha_i| \leq n^2\delta.$$

$\qquad\square$

3.2 The computational equivalence

We now provide polynomial-time reductions between HOUSINGMARKET and RAINBOWKKM, and show that the two problems have the same query complexity in the black-box model.

Theorem 3.4. *There exists a polynomial-time reduction from* HOUSINGMARKET *to* RAINBOWKKM.

Proof. Suppose $(\varepsilon, f^1, \ldots, f^n)$ is an instance of HOUSINGMARKET, and each of the n functions f^i is associated with preference set $P^i = (P^i_1, \ldots, P^i_n)$. We construct an instance $(\varepsilon', g^1, \ldots, g^n)$ of RAINBOWKKM. Set the approximation parameter to $\varepsilon' = \frac{\varepsilon}{n^2}$, and let each g^i be the covering function associated with the sparse KKM covering $C^i = (C^i_1, \ldots, C^i_n)$ of Δ_{n-1} defined by $C^i_j := \phi(P^i_j \cap \Sigma_n)$ for every $i, j \in [n]$. As ϕ is a bijection, a point $\boldsymbol{x}$ lies in C^i_j if and only

if $\phi^{-1}(\boldsymbol{x})$ lies in $P_j^i \cap \Sigma_n$. Thus, we can efficiently implement g^i as a polynomial-time function which makes at most one call to f^i.

Next, we verify that the $(C_1^i, \ldots, C_n^i)$ are sparse KKM coverings of Δ_{n-1}. Fix an agent i. It is immediate that the sets C_j^i are closed, as the P_j^i are closed and ϕ is a homeomorphism. We now check that $F_S \subseteq \bigcup_{j \in S} C_j^i$ for every $S \subseteq [n]$. Suppose that $S = [n]$. By assumption (iii) on preference sets, $P_1^i, \ldots, P_n^i$ are a covering of Σ_n and ϕ is a bijection from Σ_n to Δ_{n-1}. Hence, $C_1^i, \ldots, C_n^i$ is a covering of Δ_{n-1}. Now suppose $S \subsetneq [n]$ and let $\boldsymbol{x} \in F_S$, so $x_j = 0$ for all $j \notin S$. Let $T = [n] \setminus S$. We pick a permutation π of $[n]$ such that $T = \{\pi(1), \ldots, \pi(|T|)\}$ and $\boldsymbol{x} \in \Delta_\pi$, so $x_{\pi(j)} = 0$ for all $j \in [|T|]$. Let $\boldsymbol{p} = \phi^{-1}(\boldsymbol{x})$. By the definition of ϕ in (5), we have $p_{\pi(j)} = 1$ for all $j \in [|T|]$. It follows by assumption (ii) on preference sets that $\boldsymbol{p} \notin \bigcup_{j \in [|T|]} P_{\pi(j)}^i = \bigcup_{k \in T} P_k^i$. As $(C_1^i, \ldots, C_n^i)$ is a covering of Δ_{n-1}, we have $\phi(\boldsymbol{p}) = \boldsymbol{x} \in \bigcup_{k \in [n]} C_k^i$. It follows that $\boldsymbol{x} \in \bigcup_{k \in S} C_k^i$. Finally, we use a similar argument to show that $(C_1^i, \ldots, C_n^i)$ is sparse, i.e. $C_j^i \cap F_{[n] \setminus \{j\}} = \emptyset$ for every $i \in [n]$. Suppose $\boldsymbol{x} \in C_j^i$, so $\boldsymbol{p} = \phi^{-1}(\boldsymbol{x}) \in P_j^i$. Suppose $\boldsymbol{x} \in F_{[n] \setminus \{j\}}$, so $x_j = 0$. Letting π be a permutation with $\pi(1) = j$, we see that $p_j = 1$ and so, by assumption (ii) of preference sets we have $\boldsymbol{p} \notin P_j^i$, a contradiction to $\boldsymbol{x} \in C_j^i$.

Let $(\boldsymbol{x}, \pi)$ be an ε'-approximate Rainbow-KKM solution for our RAINBOWKKM instance. Thus there exists a point $\boldsymbol{x}^i \in C_{\pi(i)}^i$ with $\|\boldsymbol{x} - \boldsymbol{x}^i\|_1 \leq \varepsilon'$ for every $i \in [n]$. Let $\boldsymbol{p} = \phi^{-1}(\boldsymbol{x})$ and $\boldsymbol{p}^i = \phi^{-1}(\boldsymbol{x}^i)$. Firstly, $\boldsymbol{x}^i \in C_{\pi(i)}^i$ implies $\boldsymbol{p}^i \in P_{\pi(i)}^i$. Moreover, $\|\boldsymbol{p} - \boldsymbol{p}^i\|_1 = \|\phi^{-1}(\boldsymbol{x}) - \phi^{-1}(\boldsymbol{x}^i)\|_1 \leq n^2 \varepsilon' = \varepsilon$ for every $i \in [n]$, as ϕ^{-1} is n^2-Lipschitz. It follows that $(\boldsymbol{p}, \pi)$ is an ε-equilibrium solution for the initial HOUSINGMARKET instance. $\square$

Theorem 3.5. *There exists a polynomial-time reduction from* RAINBOWKKM *to* HOUSING-MARKET.

Proof. Suppose $(\varepsilon, g^1, \ldots, g^n)$ is an instance of RAINBOWKKM, and each g^i is associated with KKM covering $C_1^i, \ldots, C_n^i$ of Δ_{n-1}. Without loss of generality, we can assume that the KKM coverings are sparse (cf. Section 2.2). We construct an instance $(\varepsilon', f^1, \ldots, f^n)$ of HOUS-INGMARKET with $\varepsilon' = \frac{\varepsilon}{n}$. Each preference function f^i is associated with preference sets $P^i := (P_0^i, \ldots, P_n^i)$ given by $P_0^i = \mathbb{R}^n$ and $P_j^i = \phi^{-1}(C_j^i \cap \Delta_{n-1})$ for $j \in [n]$. As $\boldsymbol{p} \in P_0^i$ for all $\boldsymbol{p} \in \mathbb{R}^n$, and $\boldsymbol{p} \in P_j^i$ if and only if $\phi(\boldsymbol{p}) \in C_j^i \cap \Delta_{n-1}$ for any $j \in [n]$, we see that each f^i can be implemented efficiently with at most one call to g^i.

It is immediate that P^i covers $\mathbb{R}^n$, and that P_0^i is closed. Moreover, by assumption the C_j^i are closed and cover Δ_{n-1}. Hence, as ϕ^{-1} is a homeomorphism from Δ to Σ_n, we see that the P_j^i are closed and cover Σ_n. So $P^1, \ldots, P^n$ satisfy assumptions (i) and (iii) on preference sets. Finally, we argue that $p_j \geq 1$ implies $\boldsymbol{p} \notin P_j^i$ for every $j \in [n]$, in order to show that the preference sets satisfy assumption (ii). If $\boldsymbol{p} \notin \Sigma_n$, then this follows from $P_j^i \subseteq \Sigma_n$. Now suppose $\boldsymbol{p} \in \Sigma_n$ and $p_j = 1$. Then $\boldsymbol{p}$ lies in Σ_π for permutation $\pi = (j, \ldots)$ as defined in Section 3.1. Hence $\boldsymbol{x} := \phi^{-1}(\boldsymbol{p})$ satisfies $x_j = 0$ and so $\boldsymbol{x} \in F_{[n] \setminus \{j\}}$. As the KKM coverings are sparse by assumption, we get $\boldsymbol{x} \notin C_j^i$.

Now let $(\boldsymbol{p}, \pi)$ be an ε'-equilibrium solution for the HOUSINGMARKET instance. Thus there exists $\boldsymbol{p}^i \in P_{\pi(i)}^i$ with $\|\boldsymbol{p} - \boldsymbol{p}^i\|_1 \leq \varepsilon'$ for every $i \in [n]$. Note that $P_{\pi(i)}^i = \phi^{-1}(C_{\pi(i)}^i \cap \Delta_{n-1})$ is a subset of Σ_n by definition of ϕ^{-1}, so at least one entry of $\boldsymbol{p}^i$ is 0. It follows that $p_j \leq \varepsilon'$ for some $j \in [n]$. We claim that $(\boldsymbol{x}, \pi)$ with $\boldsymbol{x} = \phi(\boldsymbol{p})$ is an ε-approximate Rainbow-KKM solution for the original RAINBOWKKM instance. We have $\phi(\boldsymbol{p}^i) \in C_{\pi(i)}^i$. Moreover, as ϕ is n-Lipschitz by Lemma 3.2, $\|\phi(\boldsymbol{p}) - \phi(\boldsymbol{p}^i)\|_1 \leq n\varepsilon' = \varepsilon$. $\square$

Theorems 3.4 and 3.5 establish that HOUSINGMARKET is PPAD-complete (in the white-box model) if and only if RAINBOWKKM is. Moreover, in the proofs above we described how

to implement the covering functions using one call to the preference functions, and vice versa. Hence HOUSINGMARKET and RAINBOWKKM have the same query complexity in the black-box model.

Corollary 3.6. HOUSINGMARKET *and* RAINBOWKKM *have the same query complexity.*

It also follows from the reduction of Theorem 3.4 that computing a market equilibrium in the housing market is straightforward when the market consists of two agents. Indeed, we briefly describe how an ε-approximate Rainbow-KKM solution can be found in polynomial time using straightforward binary search on the interval Δ_1 when $n = 2$. Suppose (C_1^1, C_2^1) and (C_1^2, C_2^2) are two KKM coverings of $\Delta_1 = \text{conv}\{e^1, e^2\}$, and initialize the search interval $[\boldsymbol{x}, \boldsymbol{y}]$ with $\boldsymbol{x} = (0, 1)$ and $\boldsymbol{y} = (1, 0)$. Note that $\boldsymbol{x} \in C_1^1$ and $\boldsymbol{y} \in C_2^2$, and $\|\boldsymbol{x} - \boldsymbol{y}\|_1 = 2$. Now repeat the following until $\|\boldsymbol{x} - \boldsymbol{y}\|_1 \le \varepsilon$, at which point $\boldsymbol{x}$ and $\boldsymbol{y}$ are both ε-approximate Rainbow-KKM solutions. Compute the halfway point $\boldsymbol{z} = \frac{1}{2}\boldsymbol{x} + \frac{1}{2}\boldsymbol{y}$. If $\boldsymbol{z}$ does not lie in $C_1^1 \cup C_2^1$, then the KKM covering property guarantees $\boldsymbol{z} \in C_2^1 \cap C_1^2$, and we have found an exact Rainbow-KKM solution. Otherwise, update the left boundary $\boldsymbol{x}$ of the search interval to $\boldsymbol{z}$ if $\boldsymbol{z} \in C_1^1$ and the right boundary $\boldsymbol{y}$ of the search interval to $\boldsymbol{z}$ if $\boldsymbol{z} \in C_2^2$. As $\|\boldsymbol{x} - \boldsymbol{y}\|_1$ reduces by at least half in every iteration and we maintain the invariants $\boldsymbol{x} \in C_1^1$ and $\boldsymbol{y} \in C_2^2$ throughout, this procedure terminates correctly with a running time and query complexity of $O(\log(\frac{1}{\varepsilon}))$.

4 PPAD-Hardness of HousingMarket and RainbowKKM

We prove that HOUSINGMARKET is PPAD-hard by reducing CAKECUTTING to the intermediate problem RAINBOWKKM. Our reduction shows that any algorithm for solving the housing market with a fixed number d of agents can be applied to cutting a cake for d players. It also leads to an exponential lower bound $\Omega(\text{poly}(\frac{1}{\varepsilon}))$ on the query complexity of RAINBOWKKM and HOUSINGMARKET in the black-box model.

We also provide an alternative proof of the hardness of RAINBOWKKM by reducing 2D-SPERNER to 3D-KKM. As KKM can be reduced to RAINBOWKKM by making multiple copies of the same KKM covering, this reduction implies that HOUSINGMARKET is hard even if we have only three agents with identical preferences. CAKECUTTING and 2D-SPERNER are known to be PPAD-complete problems [Deng et al., 2012, Hollender and Rubinstein, 2023, Chen and Deng, 2009].

4.1 Reducing from the CakeCutting problem

Consider a cake represented by one-dimensional interval $[0, 1]$. The goal is to divide the cake among d players by making $d - 1$ cuts and assigning the resulting pieces to agents so that no agent envies another. The preferences of each agent $i \in [n]$ are described by a utility function u^i defined on the set of all possible pieces $[a, b] \subseteq [0, 1]$ (with $a \le b$). These utility functions satisfy

(i) $u^i(\emptyset) = 0$ and $u^i([a, b]) > 0$ for all $0 \le a < b \le 1$,

(ii) $u^i([a, b]) - u^i([a', b']) \le K(|a - a'| + |b - b'|)$ for any two intervals $[a, b], [a', b'] \subseteq [0, 1]$ and some constant K.

The first condition is known as *non-negativity* and *hungriness*, and the second condition as *Lipschitz continuity (for intervals)*.

A $(d - 1)$-cut dividing the cake $[0, 1]$ into d pieces can be represented by a vector $\boldsymbol{x} \in [0, 1]^d$ with $\sum_{i \in [n]} x_i = 1$. Entry x_k denotes the length of the k-th piece, and so the k-th piece of

the cut is the interval $I_k(\boldsymbol{x}) := [\sum_{j \in [k-1]} x_j, \sum_{j \in [k]} x_j]$. It is easy to see that the space of all $(d-1)$-cuts consists of the simplex Δ_{d-1}. A cut $\boldsymbol{x}$ is *envy-free* if there exists a one-to-one assignment of pieces to players so that each player (weakly) prefers their piece to all other pieces, i.e., there exists a permutation π of $[d]$ such that $u^i(I_{\pi(i)}(\boldsymbol{x})) \geq u^i(I_k(\boldsymbol{x}))$ for all $k \in [n]$. A cut is ε-*envy-free* if $u^i(I_{\pi(i)}(\boldsymbol{x})) \geq u^i(I_k(\boldsymbol{x})) - \varepsilon$.

Theorem 4.1 (Stromquist [1980])**.** *Under assumptions* (i) *and* (ii)*, there exists an envy-free cut.*

In the corresponding total search problem CAKECUTTING stated below, we assume that $u^1, \ldots, u^d$ are given as function oracles or polynomial-time algorithms, depending on the computational model. We let d-CAKECUTTING denote the problem with a fixed number d of players.

CakeCutting

Input: Approximation parameter $\varepsilon > 0$. Utility functions $u^1, \ldots, u^d$ for agents $[d]$ satisfying assumptions (i) and (ii).
Output: ε-envy-free cut $\boldsymbol{x} \in \Delta_{d-1}$ and corresponding assignment π of pieces to players.

Theorem 4.2 ([Hollender and Rubinstein, 2023])**.** CAKECUTTING *is PPAD-complete.*

An earlier result by Deng et al. [2012] showed that CAKECUTTING is PPAD-complete when preferences are not specified via utility functions but by means of more general 'preferences sets', analogous to our model of the housing market.

We are now ready to reduce CAKECUTTING to RAINBOWKKM, which implies that RAINBOWKKM and HOUSINGMARKET are PPAD-hard. Our proof of Theorem 4.3 relies on the representation of cake preferences in terms of utilities.

Theorem 4.3. *There exists a polynomial-time reduction from d-CAKECUTTING to d-RAINBOWKKM.*

Proof. Suppose $(\varepsilon, u^1, \ldots, u^d)$ is an instance of d-CAKECUTTING. We define a KKM covering $C^i = (C_1^i, \ldots, C_d^i)$ of Δ_{d-1} for each player $i \in [d]$ by

$$C_j^i := \{\boldsymbol{x} \in \Delta_{d-1} \mid u^i(I_j(\boldsymbol{x})) \geq u^i(I_k(\boldsymbol{x})) \text{ for all } k \in [d]\}.$$

Hence, C_j^i consists of all the cake cuts for which agent i prefers piece j. To see that C^i is indeed a KKM covering of Δ_{d-1}, consider any $S \subseteq [d]$ and let $\boldsymbol{x} \in F_S = \text{conv}\{\boldsymbol{e}^k \mid k \in S\}$. By choice of $\boldsymbol{x}$, we have $x_k = 0$ for all $k \notin S$ and $x_k > 0$ for some $k \in S$. The non-negativity property of the utility function u^i then implies that player i prefers some piece of the cake indexed by S to all the pieces not indexed by S. It follows that $\boldsymbol{x} \in C_k^i$ for some $k \in S$, so C^i is a KKM covering.

We now construct an instance of RAINBOWKKM. Fix approximation parameter $\varepsilon' = \frac{\varepsilon}{4K}$ (where K is the Lipschitz constant of the utility functions u^i), and let $g^1, \ldots, g^d$ be the KKM covering functions corresponding to $C^1, \ldots, C^d$. As $g^i(\boldsymbol{x}, j) = 1$ if $\boldsymbol{x} \in C_j^i$ and $g^i(\boldsymbol{x}, j) = 0$ otherwise, it is straightforward to see that g^i can be implemented efficiently with at most d calls to the utility function u^i of player i.

Suppose $(\boldsymbol{x}, \pi)$ is a solution to this RAINBOWKKM instance $(\varepsilon', g^1, \ldots, g^d)$. Hence, there exists a collection of points $\boldsymbol{x}^1, \ldots, \boldsymbol{x}^d$ such that $\boldsymbol{x}^i \in C_{\pi(i)}^i$ and $\|\boldsymbol{x} - \boldsymbol{x}^i\|_1 \leq \varepsilon'$ for all $i \in [d]$. We will now show that $(\boldsymbol{x}, \pi)$ forms an ε-approximate solution to the CAKECUTTING instance.

Fix a player $i \in [d]$. Our goal is to show that she ε-approximately prefers the $\pi(i)$-th piece of cut $\boldsymbol{x}$ to all other pieces specified by $\boldsymbol{x}$. Fix some $k \in [d]$, and the k-th pieces $I_k(\boldsymbol{x}) = [a, b]$ and $I_k'(\boldsymbol{x}) = [a', b']$ of cuts $\boldsymbol{x}$ and $\boldsymbol{x}^i$. Note that $a = \sum_{l=1}^{k-1} x_l$ and $a' = \sum_{l=1}^{k-1} x_l^i$, so $\|\boldsymbol{x} - \boldsymbol{x}^i\|_1 \leq \varepsilon'$

implies $|a - a'| \leq \varepsilon'$. Likewise, $|b - b'| < \varepsilon'$. By the Lipschitz continuity of the agent's utility function, we thus have

$$u^i([a, b]) - u^i([a', b']) \leq K(|a - a'| + |b - b'|) \leq 2\varepsilon' K = \frac{\varepsilon}{2}.$$

As player i prefers the $\pi(i)$-th piece of cut $\boldsymbol{x}^i$ by choice of $\boldsymbol{x}^i$, it follows that $u^i(I_{\pi(i)}(\boldsymbol{x})) \geq u^i(I_k(\boldsymbol{x}^i)) + \varepsilon$ for every $k \in [d]$. $\square$

Hollender and Rubinstein [2023] showed that CAKECUTTING with a fixed number of players $d \geq 4$ has a query complexity lower bound of $\Omega(\text{poly}(\frac{1}{\varepsilon}))$. The reduction in our proof of Theorem 4.3 thus implies the same lower bound on the query complexity for d-HOUSINGMARKET and d-RAINBOWKKM.

Corollary 4.4. *d-HOUSINGMARKET and d-RAINBOWKKM with $d \geq 4$ have a query complexity of $\Omega(\text{poly}(\frac{1}{\varepsilon}))$.*

4.2 The reduction from the 2D-Sperner problem

We now develop the reduction from 2D-SPERNER to 3-RAINBOWKKM. We do this via the intermediary problem 3D-KKM. As we know from Section 3.2 that n-HOUSINGMARKET and n-RAINBOWKKM are computationally equivalent, Theorem 2.2 then implies that both n-HOUSINGMARKET and n-RAINBOWKKM are PPAD-hard for any $n \geq 3$.

The 2D-Sperner problem. Consider the triangulation of triangle $N\Delta_2 = \text{conv}\{N\boldsymbol{e}^1, N\boldsymbol{e}^2, N\boldsymbol{e}^3\}$ with side lengths N into equilateral cells of side length 1. This is illustrated for $N = 4$ in Fig. 2. The vertices of this triangulation are $V_N := \{\boldsymbol{v} \in \mathbb{N}_0^3 \mid v_1 + v_2 + v_3 = N\}$. A *coloring* $c : V_N \to \{1, 2, 3\}$ assigns each vertex one of three colors. It is a *Sperner coloring* if $c(\boldsymbol{v}) \neq i$ when $v_i = 0$. That is, each corner $N\boldsymbol{e}^i$ of the triangle is colored in i, and each point on the boundary between $N\boldsymbol{e}^i$ and $N\boldsymbol{e}^j$ is colored in i or j. A cell of the triangulation is *trichromatic* if each of its three vertices receives a different color. Sperner's lemma (Lemma 4.5) guarantees the existence of a trichromatic cell.

Lemma 4.5 ([Sperner, 1928]). *Every Sperner coloring of every triangulation of a triangle has at least one trichromatic cell.*

This leads naturally to the computational problem of finding a trichromatic cell. We assume that the size N of the triangulation is given as an input parameter in a binary representation, as the problem is easy otherwise.

> **2D-Sperner**
>
> **Input:** The size $N \in \mathbb{N}$ of the triangle, and a Sperner coloring function $c : V_N \to \{1, 2, 3\}$ of the vertices of the triangulation of $N\Delta_2$.
> **Output:** The three vertices of a trichromatic cell in the (unit) triangulation of $N\Delta_2$.

The 2D-SPERNER problem stated above is computationally equivalent to the version considered in [Chen and Deng, 2009], which considers a triangulation of a right-angled triangle. It is straightforward to see that we can reduce between these problems. Chen and Deng [2009] settled the computational complexity of 2D-SPERNER.

Theorem 4.6 ([Chen and Deng, 2009]). *2D-SPERNER is PPAD-complete.*
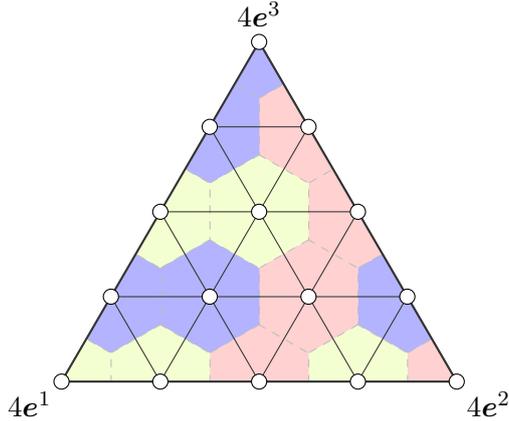
Figure 2: Triangulation of simplex Δ_2 with side length 4 into regular triangles with side length 1, drawn using solid black lines. The three corresponding covering sets C_1, C_2 and C_3 consist of the yellow, red and blue regions, respectively.

We now turn to the problem of reducing from 2D-SPERNER to KKM. Fix the triangulation of $N\Delta_2$ and a Sperner coloring c. For simplicity, consider the variant of KKM that seeks to find an ε-approximate KKM solution given a KKM covering of $N\Delta_2$, and N is given as an input parameter alongside the covering function. This is equivalent to our original definition of KKM with approximation parameter $\varepsilon' = \frac{\varepsilon}{N}$. Define the following KKM covering of $N\Delta_2$. For $i \in \{1, 2, 3\}$, let

$$C_i := \{\boldsymbol{x} \in N\Delta_2 \mid c(\boldsymbol{v}) = i \text{ for some } \boldsymbol{v} \in \operatorname{argmin}_{\boldsymbol{v} \in V} \|\boldsymbol{x} - \boldsymbol{v}\|_1\}. \tag{12}$$

Hence, C_i contains all the points whose closest vertex is i-colored. The three covering sets are illustrated in Fig. 2; each C_i is the union of identically-colored hexagons. It is straightforward to check that (C_1, C_2, C_3) is a KKM covering of $N\Delta_2$, i.e. that $F_S = \operatorname{conv}\{N\boldsymbol{e}^i \mid i \in S\}$ is contained in $\bigcup_{i \in S} C_i$ for every $S \subseteq [n]$. Indeed, fix $S \subseteq [n]$ and $\boldsymbol{x} \in F_S$, so $x_k = 0$ for all $k \notin S$. If $S = \{1\}$, then $\boldsymbol{x}$ is a vertex. As c is a Sperner coloring, we have $c(\boldsymbol{x}) = 1$ and $\boldsymbol{x} \in C_1$. If $S = \{1, 2\}$, then $\boldsymbol{x}$ lies on the boundary of $N\Delta_2$ between $N\boldsymbol{e}^1$ and $N\boldsymbol{e}^2$. The closest vertex $\boldsymbol{v} \in V$ to $\boldsymbol{x}$ also lies on this side and, as c is a Sperner coloring, we have $c(\boldsymbol{v}) \in \{1, 2\}$. Hence, $\boldsymbol{x} \in C_1 \cup C_2$. Finally, if $S = \{1, 2, 3\}$, then $\boldsymbol{x}$ lies in the interior. By construction, $\boldsymbol{x} \in C_{c(\boldsymbol{v})}$ for some vertex $\boldsymbol{v}$ closest to $\boldsymbol{x}$. So $\boldsymbol{x} \in C_1 \cup C_2 \cup C_3$.

We now establish, in Lemma 4.7, that a solution $\boldsymbol{x}$ to the instance of KKM with approximation parameter $\varepsilon = \frac{1}{8}$ and KKM covering (C^1, C^2, C^3) as defined in (12) lies in the interior of a trichromatic cell. This forms the main part of our reduction. For this result, we will need the intuitive result of Lemma 4.8 that the closest vertex to a point in a cell, or to a point 'almost in a cell', is one of the vertices of this cell. Note that we can apply a distance-preserving map to $N\Delta_2$ that maps any given cell Δ to Δ_2. Every vertex $\boldsymbol{v}$ in the resulting triangulation satisfies $\boldsymbol{v} \in \mathbb{Z}^n$ and $v_1 + v_2 + v_3 = 1$. We make use of this transformation to simplify our proofs of Lemmas 4.7 and 4.8.

Lemma 4.7. *Fix a point $\boldsymbol{x}$ in cell Δ of the triangulation. If its neighborhood $N_{1/8}(\boldsymbol{x})$ intersects with C_1, C_2 and C_3, then $\boldsymbol{x}$ lies in the interior of Δ, and Δ is trichromatic.*

Lemma 4.8. *If $\boldsymbol{x}$ lies in $N_{\frac{1}{8}}(\boldsymbol{y})$ for some point $\boldsymbol{y}$ of a cell Δ, then one of the three vertices of this cell Δ is a closest vertex to $\boldsymbol{x}$.*

Proof. Apply a distance-preserving map to $N\Delta_2$ so that $\Delta = \Delta_2$. Let $\boldsymbol{v} \notin \{\boldsymbol{e}^1, \boldsymbol{e}^2, \boldsymbol{e}^3\}$ be a vertex of the triangulation. Recall that its entries are integral and $v_1 + v_2 + v_3 = 1$. We now show that $\|\boldsymbol{x} - \boldsymbol{v}\|_1 \geq 2 - \frac{1}{8}$ and $\|\boldsymbol{x} - \boldsymbol{e}^i\| \leq \frac{4}{3}$ for some $i \in \{1, 2, 3\}$.

Suppose first that $\boldsymbol{x}$ lies on the boundary of Δ. Without loss of generality, $\boldsymbol{x}$ is a convex combination of $\boldsymbol{e}^1$ and $\boldsymbol{e}^2$ and $x_1 \geq x_2$, so $x_3 = 0$ and $x_1 + x_2 = 1$. Hence $\|\boldsymbol{x} - \boldsymbol{e}^1\|_1 = 1 - x_1 + x_2 \leq 1$, so the L_1-distance of $\boldsymbol{x}$ to one of the vertices of Δ is at most 1. Next, we show that $\|\boldsymbol{x} - \boldsymbol{v}\|_1 = |x_1 - v_1| + |x_2 - v_2| + |v_3| \geq 2$ by distinguishing between different possibilities for $\boldsymbol{v}$. If $v_1, v_2 \geq 1$, then $v_3 \leq -1$ and so $\|\boldsymbol{x} - \boldsymbol{v}\|_1 \geq 2 - x_1 - x_2 + |v_3| \geq 2$. If $v_1, v_2 \leq 0$, with at least one strict inequality, then $v_3 \geq 2$ and so $\|\boldsymbol{x} - \boldsymbol{v}\|_1 = x_1 + x_2 - v_1 - v_2 + |v_3| \geq 4$. If $v_1 = 0$ and $v_2 \geq 2$, then $v_3 \leq -1$ and so $\|\boldsymbol{x} - \boldsymbol{v}\|_1 = x_1 + x_2 + v_2 + |v_3| \geq 2$. If $v_1 < 0$ and $v_2 = 1$, then $v_3 \geq 1$ implies $\|\boldsymbol{x} - \boldsymbol{v}\|_1 \geq 2$. Finally, if $v_1 < 0$ and $v_2 > 1$, then $v_2 - v_1 \geq 3$ guarantees the same.

Now suppose that $\boldsymbol{x}$ lies in the interior of Δ. Without loss of generality, let $x_1 \geq x_2 \geq x_3$, so $\|\boldsymbol{x} - \boldsymbol{e}^1\|_1 \leq \|\boldsymbol{x} - \boldsymbol{e}^2\|_1$ and $\|\boldsymbol{x} - \boldsymbol{e}^1\|_1 \leq \|\boldsymbol{x} - \boldsymbol{e}^3\|_1$. Moreover, $\|\boldsymbol{x} - \boldsymbol{e}^1\|_1 = 1 - x_1 + x_2 + x_3$ is maximized when $\boldsymbol{x}$ is the barycenter of Δ_2, so $x_1 = x_2 = x_3 = \frac{1}{3}$. Hence $\|\boldsymbol{x} - \boldsymbol{e}^i\|_1 \leq \frac{4}{3}$. Let $\boldsymbol{y}$ be the point at which the line segment from $\boldsymbol{x}$ to $\boldsymbol{v}$ intersects the boundary of Δ. It is straightforward that $\|\boldsymbol{x} - \boldsymbol{v}\|_1 = \|\boldsymbol{x} - \boldsymbol{y}\|_1 + \|\boldsymbol{y} - \boldsymbol{v}\|_1 \geq \|\boldsymbol{y} - \boldsymbol{v}\|_1 \geq 2$, where the second inequality follows from the previous paragraph.

Finally, suppose $\boldsymbol{x}$ does not lie in Δ. By construction, there exists $\boldsymbol{y}$ on the boundary of Δ with $\|\boldsymbol{x} - \boldsymbol{y}\|_1 \leq \frac{1}{8}$. By the triangle inequality and the previous paragraph, the distance from $\boldsymbol{x}$ to $\boldsymbol{e}^1$, $\boldsymbol{e}^2$ or $\boldsymbol{e}^3$ is at most $1 + \frac{1}{8}$, and the distance from $\boldsymbol{x}$ to $\boldsymbol{y}$ is at least $2 - \frac{1}{8}$, so the result holds. $\qquad\square$

Proof of Lemma 4.7. Without loss of generality, let Δ be the unit simplex Δ_2 (by applying a suitable distance-preserving map) and suppose each of its vertices $\boldsymbol{e}^i$ is i-colored. Let $\boldsymbol{z} := (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ be the barycenter of Δ.

Assume that $N_{1/8}(\boldsymbol{x}) \cap C_i \neq \emptyset$, so there exist $\boldsymbol{x}^i \in C_i$, for all $i \in \{1, 2, 3\}$. Suppose we can show that $\boldsymbol{x}$ lies in $N_{1/2}(\boldsymbol{z})$. Then $\boldsymbol{x}^i \in N_{5/8}(\boldsymbol{z})$ for every i. As $N_{5/8}(\boldsymbol{z})$ is contained in Δ, the points $\boldsymbol{x}, \boldsymbol{x}^1, \boldsymbol{x}^2$ and $\boldsymbol{x}^3$ then all lie in the interior of Δ. But as $\boldsymbol{x}^i \in C_i$, and $\boldsymbol{e}^1, \boldsymbol{e}^2$ and $\boldsymbol{e}^3$ are the three closest vertices to any point in Δ by Lemma 4.8, it then follows from the definition of C^i that Δ is trichromatic.

It remains to show that $\boldsymbol{x} \in N_{1/2}(\boldsymbol{z})$. Suppose, for the sake of contradiction, that $\boldsymbol{x} \notin N_{1/2}(\boldsymbol{z})$. This means that $\|\boldsymbol{x} - \boldsymbol{z}\|_1 = |x_1 - \frac{1}{3}| + |x_2 - \frac{1}{3}| + |x_3 - \frac{1}{3}| > \frac{1}{2}$. Without loss of generality, assume that $x_1 \geq x_2 \geq x_3$. As $0 \leq \boldsymbol{x} \neq \boldsymbol{z}$ and $x_1 + x_2 + x_3 = 1$, we have $x_1 > x_3$. Below, we will argue that

$$x_1^3 > x_3^3 \text{ or } x_2^3 > x_3^3. \tag{13}$$

This implies $\|\boldsymbol{x}^3 - \boldsymbol{e}^3\|_1 - \|\boldsymbol{x}^3 - \boldsymbol{e}^1\|_1 = x_1^3 - x_3^3 > 0$ or $\|\boldsymbol{x}^3 - \boldsymbol{e}^3\|_1 - \|\boldsymbol{x}^3 - \boldsymbol{e}^2\|_1 = x_2^3 - x_3^3 > 0$. The fact that $\boldsymbol{x}^3$ is thus closer to $\boldsymbol{e}^1$ or $\boldsymbol{e}^2$ leads to the contradiction $\boldsymbol{x}^3 \notin C^3$.

We now argue that (13) holds. As $\|\boldsymbol{x} - \boldsymbol{x}^3\|_1 \leq \frac{1}{8}$, we have $|x_i - x_i^3| \leq \frac{1}{8}$ for every i. Hence (13) holds immediately if $x_1 - x_3 > \frac{1}{4}$ or $x_2 - x_3 > \frac{1}{4}$. It remains to prove this last claim. As $\boldsymbol{x}$ lies in Δ, we have $x_1 + x_2 + x_3 = 1$. Hence, $x_1 \geq x_2 \geq x_3$ and $x_1 > x_3$ implies $x_1 > \frac{1}{3}$ and $x_3 < \frac{1}{3}$. It follows that $\|\boldsymbol{x} - \boldsymbol{z}\|_1 = x_1 - x_3 + |x_2 - \frac{1}{3}| > \frac{1}{2}$, so $x_1 - x_3 > \frac{1}{4}$ or $|x_2 - \frac{1}{3}| > \frac{1}{4}$. If $x_1 - x_3 > \frac{1}{4}$, we are done. So suppose $x_1 - x_3 \leq \frac{1}{4}$ and $|x_2 - \frac{1}{3}| > \frac{1}{4}$. This implies $x_2 > \frac{1}{4} + \frac{1}{3}$ or $x_2 < \frac{1}{3} - \frac{1}{4}$. Suppose the latter holds, so $x_2 < \frac{1}{3} - \frac{1}{4}$. We also bound x_1 from above by noting that $x_1 - x_3 \leq \frac{1}{4}$ and $x_3 < \frac{1}{3}$ imply $x_1 < \frac{1}{4} + \frac{1}{3}$. It follows that $x_1 + x_2 + x_3 < \frac{1}{4} + \frac{1}{3} + \frac{1}{3} - \frac{1}{4} + \frac{1}{3} = 1$, a contradiction. Hence, we must have $x_2 > \frac{1}{4} + \frac{1}{3}$. It follows that $x_2 - x_3 > x_2 - \frac{1}{3} > \frac{1}{4}$, and we are done. $\qquad\square$

Theorem 4.9. *There exists a polynomial-time reduction from* 2D-SPERNER *to 3D-KKM.*

Proof. Suppose (N, c) is a 2D-SPERNER instance. We construct an instance (ε, g) of KKM with $\varepsilon = \frac{1}{8}$. The covering function g is associated with the KKM covering C_1, C_2, C_3 of $N\Delta_2$ constructed from u as described in (12). Note that for any point $\boldsymbol{x} \in N\Delta_2$, we have $g(\boldsymbol{x}, i) = 1$ if a nearest vertex to $\boldsymbol{x}$ is i-colored, and $g(\boldsymbol{x}, i) = 0$ otherwise. Given an arbitrary point $\boldsymbol{x} \in N\Delta_2$, we can round up or down the entries of $\boldsymbol{x}$ to determine its nearest vertices $v^1 = (\lfloor x_1 \rfloor, \lceil x_2 \rceil, \lceil x_3 \rceil)$, $v^1 = (\lceil x_1 \rceil, \lfloor x_2 \rfloor, \lceil x_3 \rceil)$ and $v^3 = (\lceil x_1 \rceil, \lceil x_2 \rceil, \lfloor x_3 \rfloor)$. (If $\boldsymbol{x}$ is a vertex of the triangulation, v_1, v_2 and v_3 will coincide; if $\boldsymbol{x}$ lies on a face separating two cells, two of these vertices will coincide.) Then, with at most 3 calls to the coloring function, we can determine whether v^1, v^2 or v^3 are i-colored. This implements the covering function g efficiently.

Now suppose $\boldsymbol{x}$ is a solution to the KKM instance (ε, g), so there exists a point $\boldsymbol{x}^i \in C_i$ with $\|\boldsymbol{x}^i - \boldsymbol{x}\|_1 \leq \frac{1}{8}$ for every $i \in \{1, 2, 3\}$. The point $\boldsymbol{x}$ lies in some cell Δ of the triangulation. By Lemma 4.7, $\boldsymbol{x}$ lies in the interior of Δ, and the cell is trichromatic. As above, we can compute its vertices v^1, v^2 and v^3 efficiently by rounding up and down to recover a solution to the 2D-SPERNER instance (N, c). $\qquad\square$

Theorem 4.10. *There exists a polynomial-time reduction from 3D-KKM to 3-RAINBOWKKM.*

Proof. Suppose (ε, N, g) is an instance of KKM associated with KKM covering $C = (C_1, C_2, C_3)$ of $N\Delta_2$. We reduce to an instance $(\varepsilon', h^1, h^2, h^3)$ of 3-RAINBOWKKM with approximation parameter $\varepsilon' := \frac{\varepsilon}{N}$ and three identical copies of the KKM covering (D_1, D_2, D_3) of Δ_2 defined by $D_i = \frac{1}{N}C_i$. A ε'-approximate solution $\boldsymbol{x}$ to the latter, together with any permutation π of $[3]$, is an ε-approximate solution to the former, and the covering functions h^i for the 3-RAINBOWKKM instance can be constructed efficiently. $\qquad\square$

We can now combine the equivalence of n-HOUSINGMARKET and n-RAINBOWKKM shown in Section 3.2 with Theorems 2.2 and 4.10 to get the following hardness result.

Corollary 4.11. n-HOUSINGMARKET *and* n-RAINBOWKKM *are PPAD-hard for any* $n \geq 3$.

5 Membership of HousingMarket and RainbowKKM in PPAD

We now show that RAINBOWKKM lies in PPAD by reducing it to SPERNER. By Theorem 3.4, this also immediately implies that HOUSINGMARKET lies in PPAD. Our reduction uses the same labeling and coloring technique employed by [Deng et al., 2012] to reduce CAKECUTTING to SPERNER.

Suppose $C^1, \ldots, C^n$ is a family of KKM coverings. Like Deng et al. [2012], we first we define the *Kuhn triangulation* of the 'large' cube $[0, N]^{n-1}$ obtained by dividing the cube into N^{n-1} unit cubes, and then subdividing each unit cube into $(n-1)!$ simplicial cells. Next, we label *and* color each vertex of this triangulation with one of colors $[n-1]_0$ and labels $[n-1]_0$.[2] In order to define our labeling rule and coloring function, we first divide the 'large' cube into $(n-1)!$ 'large' simplices. Then we associate each vertex $\boldsymbol{v}$ with its barycentric coordinates $\alpha(\boldsymbol{v})$ w.r.t. the vertices of the large simplex in which $\boldsymbol{v}$ lies. The labeling rule and coloring function are then specified with respect to $\alpha(\boldsymbol{v})$. We use the same labeling rule as [Deng et al., 2012]; it guarantees that the vertices of a cell all receive different labels. Our coloring function then assigns each vertex $\boldsymbol{v}$, labeled i, the smallest color j for which $\alpha(\boldsymbol{v})$ lies in covering set C_j^i. Sperner's lemma implies that there exists a panchromatic cell. Thus for every label i, the i-labeled vertex of this cell is colored differently. Finally, as the L_1-distance between the barycentric coordinates of any two vertices in a cell of the Kuhn triangulation is at most $\frac{n}{N}$, we arrive at an $\frac{n}{N}$-approximate Rainbow-KKM solution to the RAINBOWKKM problem with KKM coverings $C^1, \ldots, C^n$.

---

[2]Unlike Deng et al. [2012], we label and color the entire triangulated large cube, and not just a subdivision of it. This allows us to reduce our problem to the standard SPERNER problem defined on the subdivision of a cube.
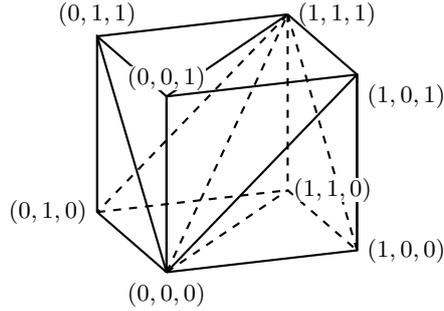
Figure 3: The Kuhn triangulation of the unit cube.

5.1 Kuhn's triangulation

We describe Kuhn's method for triangulating a cube introduced by Scarf [1982] and Deng et al. [2012]. We begin by defining Kuhn's triangulation of the n-dimensional unit cube $[0,1]^n$. For each permutation π of $[n]$, define the base simplex $\hat{\Delta}_\pi := \{ \boldsymbol{x} \in [0,1]^n \mid x_{\pi(1)} \geq x_{\pi(2)} \geq \cdots \geq x_{\pi(n)} \}$. The vertices of this simplex are given by $\boldsymbol{v}_\pi^0 := \boldsymbol{0}$ and $\boldsymbol{v}_\pi^i := \boldsymbol{v}_\pi^{i-1} + \boldsymbol{e}^{\pi(i)}$ for all $i \in [n]$. It is straightforward that the unit cube consists of the union of $\hat{\Delta}_\pi$ taken over all permutations π, and that the interiors of the base simplices do not overlap. The Kuhn triangulation of the unit cube is illustrated for $n = 3$ in Fig. 3.

Next, fix a positive integer $N \in \mathbb{N}$. We describe a way to triangulate the cube $[0,N]^n$ of side length N. We also call this the Kuhn triangulation of the "large cube". First, divide the cube into N^d unit cubes. Then triangulate each unit cube into *base simplices* (or *cells*) by applying Kuhn's triangulation as described in the previous paragraph. Hence each unit cube contains $n!$ cells, one for each permutation π. In particular, if $\boldsymbol{x}$ is the smallest point in the cube, then its cell associated with permutation π is $\boldsymbol{x} + \hat{\Delta}_\pi$, and is the convex hull of vertices $\boldsymbol{v}_\pi^0 := \boldsymbol{x}$ and $\boldsymbol{v}_\pi^i := \boldsymbol{v}_\pi^{i-1} + \boldsymbol{e}^{\pi(i)}$ for all $i \in [n]$. The set of all vertices of the triangulation is $V_N := [n]_0^n$.

Alternatively, we can also divide the large cube $[0,N]^n$ into simplices $N\hat{\Delta}_\pi := \{ N\boldsymbol{x} \mid \boldsymbol{x} \in \hat{\Delta}_\pi \}$ for all π. (This is analogous to performing Kuhn's triangulation of the unit cube, scaled up by a factor of N.) We will refer to the $n!$ resulting simplices $N\hat{\Delta}_\pi$ as the "large simplices". Kuhn's triangulation of the large cube has the important property, stated in Lemma 5.1, that every cell is contained entirely in one of the large simplices. For more details on Kuhn's triangulation, see [Deng et al., 2012] and [Scarf, 1982].

Lemma 5.1 ([Deng et al., 2012]). *Every cell of Kuhn's triangulation of the large cube $[0,N]^n$ is contained in a large simplex $N\hat{\Delta}_\pi$ for some permutation π of $[n]$.*

For every vertex $\boldsymbol{v} \in V_N$, let $\alpha(\boldsymbol{v})$ denote its barycentric coordinates $\alpha(\boldsymbol{v}) = (\alpha(\boldsymbol{v})_0, \ldots, \alpha(\boldsymbol{v})_n)$ w.r.t. the vertices $v_\pi^0, \ldots, v_\pi^n$ of the large simplex $N\hat{\Delta}_\pi$ in which $\boldsymbol{v}$ lies. Note that $\alpha(\boldsymbol{v})$ is $(n+1)$-dimensional and satisfies $\alpha(\boldsymbol{v}) \geq \boldsymbol{0}$ and $\sum_{i \in [n]_0} \alpha(\boldsymbol{v})_i = 1$. If a vertex lies in two large simplices $N\hat{\Delta}_\pi$ and $N\hat{\Delta}_\tau$, then Lemma 5.2 states that its barycentric coordinates w.r.t. the vertices of $N\hat{\Delta}_\pi$ and $N\hat{\Delta}_\tau$ are identical. This ensures that $\alpha(\boldsymbol{v})$ is well-defined. Moreover, Lemma 5.3 shows that the L_1-distance between the barycentric coordinates of two vertices of a cell is at most $\frac{n+1}{N}$. The proof of Lemma 5.3 is essentially identical to the proof in [Deng et al., 2012], and is stated for completeness.

Lemma 5.2. *Suppose $\boldsymbol{x}$ lies in $N\hat{\Delta}_\pi$ and $N\hat{\Delta}_\pi$. Then*

$$ \sum_{i=0}^n \alpha_i v_\pi^i = \sum_{i=0}^n \alpha_i v_\tau^i $$

20

for some barycentric coordinate vector $\alpha \in \Delta_n$.

Proof. We will construct a sequence $\pi = \pi^0, \pi^1, \ldots, \pi^n = \tau$ of permutations of $[n]$ that satisfies three properties. Let $\boldsymbol{\alpha}^k$ be the barycentric coordinates of $\boldsymbol{x}$ w.r.t. the vertices $v^0 = \boldsymbol{0}$ and $v^i_{\pi^k} = v^{i-1}_{\pi^k} + \boldsymbol{e}^{\pi(i)}$ ($\forall i \in [n]$) of $N\hat{\Delta}_{\pi^k}$. The three properties are:

(i) $\pi^k_1, \ldots, \pi^k_k = \tau_1, \ldots, \tau_k$ for all $k \in [n]$,

(ii) $\boldsymbol{x} \in N\hat{\Delta}_{\pi^k}$ for all $k \in [n]_0$,

(iii) $\boldsymbol{\alpha}^k = \boldsymbol{\alpha}^{k-1}$ for all $k \in [n]$.

These three properties immediately imply the lemma.

We now define the sequence of permutations. Let $\pi^0 := \pi$. We then iteratively define π^k from π^{k-1} as follows. Intuitively, the k-th iteration moves the element $\tau(k)$ in permutation π^{k-1} to the k-th position of π^k. Formally, we let $j \in [n]$ be the index for which $\pi^k(j) = \tau(k)$ and define

$$\pi^k(l) := \begin{cases} \pi^{k-1}(j) & \text{if } l = k, \\ \pi^{k-1}(l-1) & \text{if } l \in \{k+1, \ldots, j\}, \\ \pi^{k-1}(l) & \text{else} \end{cases}$$

It is immediate that property (i) holds, and that $j \geq k$ in every iteration; so the sequence of permutations is well-defined. We now prove property (ii) by induction on k. The base case $\boldsymbol{x} \in \hat{\Delta}_{\pi^0}$ holds by definition. Now suppose that $\boldsymbol{x} \in N\hat{\Delta}_{\pi^{k-1}}$, and let j be as chosen above to define π^k. Hence, $x_{\pi^{k-1}(k)} \geq \cdots \geq x_{\pi^{k-1}(j)}$. Moreover, $x \in N\hat{\Delta}_\tau$ implies $x_{\pi^{k-1}(j)} \geq x_{\pi^{k-1}(k)}$ by our choice of j. It follows that $x_{\pi^{k-1}(k)} = \cdots = x_{\pi^{k-1}(j)}$, so $\boldsymbol{x} \in N\hat{\Delta}_{\pi^k}$. Finally, we prove property (iii). For this, note that the vertices corresponding to π^{k-1} and π^k satisfy $v^i_{\pi^{k-1}} = v^i_{\pi^k}$ for $i \in \{1, \ldots, k-1\}$ and $i \in \{j, \ldots, n\}$ by construction of the vertices of $N\hat{\Delta}_{\pi^{k-1}}$ and $N\hat{\Delta}_{\pi^k}$. Moreover, $\boldsymbol{x} \in N\hat{\Delta}_{\pi^{k-1}}$ and $\boldsymbol{x} \in N\hat{\Delta}_{\pi^k}$ inductively imply $\alpha^{k-1}_i = \alpha^k_i$ for $i \in \{j, \ldots, n\}$, then $\alpha^{k-1}_i = 0 = \alpha^k_i$ for $i \in \{k, \ldots, i-1\}$, and finally $\alpha^{k-1}_i = \alpha^k_i$ for $i \in \{1, \ldots, k-1\}$. $\square$

Lemma 5.3 (cf. [Deng et al., 2012]). *For any two vertices $\boldsymbol{y}$ and $\boldsymbol{z}$ of a cell, $\|\alpha(\boldsymbol{y}) - \alpha(\boldsymbol{z})\|_1 \leq \frac{n+1}{N}$.*

Proof. Let $\boldsymbol{y}$ and $\boldsymbol{z}$ be two vertices of a cell Δ that is contained in large simplex $N\hat{\Delta}_\pi$. Without loss of generality, we assume that $\pi = (1, \ldots, n)$. (For other permutations, the result holds by symmetry). The vertices of the large simplex are $\boldsymbol{v}^0 = \boldsymbol{0}$ and $\boldsymbol{v}^i = \boldsymbol{v}^{i-1} + \boldsymbol{e}^i$ for every $i \in [n]$.

Let $\boldsymbol{x}$ be the smallest point in the unit cube containing Δ, and the $n+1$ vertices of Δ be $\boldsymbol{w}^0_\tau = \boldsymbol{x}$ and $\boldsymbol{w}^i_\tau = \boldsymbol{w}^{i-1}_\tau + \boldsymbol{e}^{\tau(i)}$ ($\forall i \in [n]$) for some permutation τ. Without loss of generality, we let $\boldsymbol{y} = \boldsymbol{w}^k$ and $\boldsymbol{z} = \boldsymbol{w}^l$ with $k < l$. By construction of the vertices of the large simplex, the barycentric coordinates of $\boldsymbol{w}^i_\tau$ for $i \in [n]$ are given by $\alpha(\boldsymbol{w}^i_\tau) = \alpha(\boldsymbol{w}^{i-1}_\tau) + \frac{1}{N}\boldsymbol{e}^{\tau(i)} - \frac{1}{N}\boldsymbol{e}^{\tau(i)-1}$. Hence,

$$\alpha(\boldsymbol{z}) - \alpha(\boldsymbol{y}) = \frac{1}{N}\sum_{i=k+1}^{l}\left(\boldsymbol{e}^{\tau(i)} - \boldsymbol{e}^{\tau(i)-1}\right) \in \frac{1}{N}\{-1, 0, 1\}^{n+1}.$$

$\square$

Starting from vertex $\boldsymbol{0}$ and using the same argument as in the proof of Lemma 5.3, we can iteratively show for every vertex $\boldsymbol{v} \in V_N$ that the entries of $\alpha(\boldsymbol{v})$ can be expressed as a proper fraction with denominator N.

Corollary 5.4. *For any vertex $\boldsymbol{v} \in V_N$ and $i \in [n]_0$, we have $\alpha(\boldsymbol{v})_i \in \left\{\frac{1}{N}, \frac{2}{N}, \ldots, 1\right\}$.*

Finally, we use the mapping α of vertices to their barycentric coordinates to define a labeling of the vertices. We use the same labeling rule as Deng et al. [2012], but extend it to the entire triangulated large cube. Lemma 5.5 shows that, under this labeling, every cell's vertices are labeled with all numbers $[n]_0$. (We restate the proof from [Deng et al., 2012] for completeness.) Such a labeling is called a *Simmons-Su labeling*. Note also that $\mathcal{L}(\boldsymbol{v})$ is well-defined, as $N\alpha(\boldsymbol{v})_i$ is integral by Corollary 5.4. For each $\boldsymbol{v} \in V$, let

$$\mathcal{L}(\boldsymbol{v}) := \sum_{i \in [n]_0} iN\alpha(\boldsymbol{v})_i \mod n+1. \tag{14}$$

Lemma 5.5 (cf. [Deng et al., 2012]). *The labeling $\mathcal{L}$ is a Simmons-Su labeling.*

Proof. Suppose Δ is a cell in the triangulation of the large cube. As in the proof of Lemma 5.3, we see that the barycentric coordinates of the vertices $\boldsymbol{v}_\pi^0, \ldots, \boldsymbol{v}_\pi^n$ of Δ satisfy $\alpha(\boldsymbol{v}_\pi^i) = \alpha(\boldsymbol{v}_\pi^{i-1}) + \frac{1}{N}\boldsymbol{e}^{\pi(i)} - \frac{1}{N}\boldsymbol{e}^{\pi(i)-1}$. Hence every vertex of Δ receives a different label, as

$$\mathcal{L}(\boldsymbol{v}_\pi^i) \equiv \mathcal{L}(\boldsymbol{v}_\pi^{i-1}) + \pi(i) - (\pi(i) - 1) \equiv \mathcal{L}(\boldsymbol{v}_\pi^{i-1}) + 1 \mod n+1.$$

$\square$

5.2 The Sperner problem

We now define the SPERNER problem for arbitrary dimensions. In contrast to 2D-SPERNER, which was defined using a triangulated triangle, Papadimitriou [1994] introduces the SPERNER problem for arbitrary dimensions using a 'large' cube divided into unit cubes. Suppose, for instance, that we are given a (Kuhn) triangulation of the large cube $[0, N]^n$. A coloring $c : V_N \to [n]_0$ of its vertices is a *Sperner coloring* if it satisfies

1. $c(\boldsymbol{v}) \neq i$ if $v_i = 0$, for all $i \in [n]$,

2. $c(\boldsymbol{v}) \neq 0$ if $v_i = N$ for any $i \in [n]$.

By embedding the triangulated cube into the triangulation of an n-dimensional simplex, it can be shown that Sperner's lemma implies the existence of a *panchromatic* cell. That is, a cell whose $n + 1$ vertices all receive a different color. This inspires the following total search problem, analogous to 2D-SPERNER. Papadimitriou [1994, Proof of Theorem 3] shows that SPERNER lies in PPAD.

Sperner

Input: Dimension $n \in \mathbb{N}$ and side length $N \in \mathbb{N}$ specifying Kuhn triangulation of large cube $[0, N]^n$. Sperner coloring function $c : [N]_0^n \to [n]_0$.
Output: The vertices of a panchromatic cell.

Theorem 5.6 ([Papadimitriou, 1994]). SPERNER *is in PPAD.*

5.3 The reduction from RainbowKKM to Sperner

Suppose $C^1, \ldots, C^n$ is a family of n KKM coverings of Δ_n and let $\varepsilon > 0$ be an approximation parameter. For notational convenience, index the sets of each covering C^i by $[n-1]_0$ instead of $[n]$, so $C^i = (C_0^i, C_1^i, \ldots, C_{n-1}^i)$. Without loss of generality, we assume that the KKM coverings are sparse (cf. Section 2.2). Consider the Kuhn triangulation of the cube $[0, N]^{n-1}$ with $N = \frac{n}{\varepsilon}$ and let its vertices be the set V. Note that the cube is $(n-1)$-dimensional. The vertices of this

triangulation are labeled with labels $[n-1]_0$ according to labeling $\mathcal{L}$. The key to our reduction is to define the coloring $c : V \to [n-1]_0$ as

$$c(\boldsymbol{v}) = \min \left\{ j \in [n-1]_0 \mid \alpha(\boldsymbol{v}) \in C_j^{\mathcal{L}(\boldsymbol{v})} \right\}. \tag{15}$$

We now argue that c is a Sperner coloring by verifying both conditions (1) and (2) of the definition in Section 5.2. Fix some vertex $\boldsymbol{v} \in V$. If $v_j = 0$ for some $j \in [n-1]$, then $\alpha_j = 0$, and so $\alpha \in F_{[n-1]_0 \setminus \{j\}}$. As the KKM coverings are sparse, we have $\alpha \notin C_j^i$ for all coverings C^i. This verifies condition (1). Now suppose $v_j = N$ for some $j \in [n-1]$. This implies $\sum_{k=1}^{n-1} \alpha_k = 1$, and so $\alpha_0 = 0$. By the same argument as above, we get $\alpha \notin C_0^i$ for all coverings C^i. This verifies condition (2).

Hence there exists a cell of the triangulation that is panchromatic, so every vertex of the cell has a different color. Fix such a cell. As $\mathcal{L}$ is a Simmons-Su labeling, each vertex of the cell also has a different label. Label these vertices $\boldsymbol{v}^0, \ldots, \boldsymbol{v}^{n-1}$ so that $\mathcal{L}(\boldsymbol{v}^i) = i$. Define the permutation π of $[n-1]_0$ as $\pi(i) = c(\boldsymbol{v}^i)$. Then $\alpha(\boldsymbol{v}^i) \in C_{\pi(i)}^i$. By Lemma 5.3, the L_1-distance between any two $\boldsymbol{v}^i$ and $\boldsymbol{v}^j$ is at most $\frac{n}{N} = \varepsilon$. It thus follows that any $\boldsymbol{v}^i$ together with π form an ε-approximate Rainbow-KKM solution for the family of KKM coverings $C^1, \ldots, C^n$.

Theorem 5.7. *There is a polynomial-time reduction from* RAINBOWKKM *to* SPERNER.

Proof. Suppose $(\varepsilon, g^1, \ldots, g^n)$ is an instance of RAINBOWKKM, and each g^i is associated with KKM covering C^i. It remains to argue that we can efficiently construct a routine that implements coloring function c using functions g^i. Recall the definition of the coloring c from (15). In order to determine the color of vertex $\boldsymbol{v}$, we first compute its barycentric coordinates α. Then compute its label $l = \sum_{i \in [n-1]_0} iN\alpha(\boldsymbol{v}) \mod n$. Finally, determine the smallest $j \in [n-1]_0$ for which $\alpha \in C_j^l$, using at most n calls to the covering function g^l. $\square$

In the white-box model, Theorems 5.6 and 5.7 imply that HOUSINGMARKET and RAINBOWKKM lie in PPAD. Moreover, the reduction also immediately implies a brute-force approach for RAINBOWKKM (and thus also for HOUSINGMARKET): compute the color $c(\boldsymbol{v})$ of each vertex $\boldsymbol{v} \in V$ and check for each cell whether it is panchromatic. In the black-box model, this algorithm has an exponential query complexity $O(N^n)$ in the input size $\log \frac{1}{\varepsilon}$ of the approximation parameter ε. We can improve this upper bound on the query complexity of HOUSINGMARKET and RAINBOWKKM. [Deng et al., 2011] show that SPERNER has a query complexity of $\Theta(N^{n-1})$ [Deng et al., 2011]. Our reduction in the proof of Theorem 5.7 solves an instance of RAINBOWKKM with $O(N^{n-1})$ queries to the coloring function c and every such query costs at most n queries to the covering function.

Corollary 5.8. HOUSINGMARKET *and* RAINBOWKKM *have query complexity* $O(N^{n-1})$.

6   Conclusion

We have shown that HOUSINGMARKET and the closely related RAINBOWKKM are PPAD-complete. Whereas the housing market is tractable with two agents, we show that PPAD-completeness holds even when we restrict the market to $n \geq 3$ agents with identical preferences. Moreover, we describe exponential upper and lower bounds for the query complexity of solving the two problems in the white-box model that hold when $n \geq 4$.

While our results may be perceived as negative for market design, it is not uncommon for PPAD-complete problems to be implemented for allocation in practice (e.g., approximate competitive equilibrium with equal incomes for course allocation). Rather we hope that our

results will stimulate further examination of the computational complexity of allocation systems with income effects. Three sets of open questions naturally arise.

First, what is the computational complexity of HOUSINGMARKET when preferences are provided as utility functions instead of preference sets? Does hardness continue to hold under the natural assumptions of monotonicity in money and no externalities as in Quinzii [1984] and Svensson [1984]? What is the complexity of computing competitive equilibrium in model with multi-good demand under assumptions, such as net substitutability condition, that guarantee existence [Baldwin et al., 2023]? Secondly, what is the exact mapping between the assumptions in HOUSINGMARKET and the monotonicity assumptions on valuations in CAKECUTTING imposed by Deng et al. [2012] and Hollender and Rubinstein [2023] to achieve polynomial-time algorithms for cutting a cake with 3 and 4 players? Thirdly, what algorithms might work well in practice for approximating equilibria in the presence of income effects?

References

B. Baisa. Auction design without quasilinear preferences. *Theoretical Economics*, 12(1):53–78, 2017.

E. Baldwin, R. Jagadeesan, P. Klemperer, and A. Teytelboym. The equilibrium existence duality. *Journal of Political Economy*, 131(6):000–000, 2023.

J.-P. Benoit and V. Krishna. Multiple-object auctions with budget constrained bidders. *The Review of Economic Studies*, 68(1):155–179, 2001.

S. Bhattacharya, G. Goel, S. Gollapudi, and K. Munagala. Budget constrained auctions with heterogeneous items. In *Proceedings of the Forty-Second ACM Symposium on Theory of Computing*, STOC '10, page 379–388, New York, NY, USA, 2010. Association for Computing Machinery. doi: 10.1145/1806689.1806743.

Y.-K. Che and I. Gale. Standard auctions with financially constrained bidders. *The Review of Economic Studies*, 65(1):1–21, 1998.

T. Chen, X. Chen, B. Peng, and M. Yannakakis. Computational hardness of the Hylland-Zeckhauser scheme. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2253–2268. SIAM, 2022.

X. Chen and X. Deng. On the complexity of 2D discrete fixed point problem. *Theoretical Computer Science*, 410(44):4448–4456, 2009.

X. Chen and S.-H. Teng. Spending is not easier than trading: on the computational equivalence of Fisher and Arrow-Debreu equilibria. In *International Symposium on Algorithms and Computation*, pages 647–656. Springer, 2009.

X. Chen, D. Dai, Y. Du, and S.-H. Teng. Settling the complexity of Arrow-Debreu equilibria in markets with additively separable utilities. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 273–282. IEEE, 2009a.

X. Chen, X. Deng, and S.-H. Teng. Settling the complexity of computing two-player nash equilibria. *Journal of the ACM (JACM)*, 56(3):1–57, 2009b.

X. Chen, D. Paparas, and M. Yannakakis. The complexity of non-monotone markets. *Journal of the ACM (JACM)*, 64(3):1–56, 2017.

C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a Nash equilibrium. *Communications of the ACM*, 52(2):89–97, 2009.

X. Deng, Q. Qi, A. Saberi, and J. Zhang. Discrete fixed points: Models, complexities, and applications. *Mathematics of Operations Research*, 36(4):636–652, 2011.

X. Deng, Q. Qi, and A. Saberi. Algorithmic solutions for envy-free cake cutting. *Operations Research*, 60(6):1461–1476, 2012.

S. Dobzinski, R. Lavi, and N. Nisan. Multi-unit auctions with budget limits. *Games and Economic Behavior*, 74(2):486–503, 2012.

D. Gale. Equilibrium in a discrete exchange economy with money. *International Journal of Game Theory*, 13(1):61–64, 1984.

P. W. Goldberg. A survey of PPAD-completeness for computing Nash equilibria. *arXiv preprint arXiv:1103.2709*, 2011.

A. Hollender and A. Rubinstein. Envy-free cake-cutting for four agents. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 113–122. IEEE, 2023.

A. Hylland and R. Zeckhauser. The efficient allocation of individuals to positions. *Journal of Political economy*, 87(2):293–314, 1979.

B. Knaster, K. Kuratowski, and S. Mazurkiewicz. Ein Beweis des Fixpunktsatzes für n-dimensionale Simplexe. *Fundamenta Mathematicae*, 14(1):132–137, 1929.

T. C. Koopmans and M. Beckmann. Assignment problems and the location of economic activities. *Econometrica: journal of the Econometric Society*, pages 53–76, 1957.

S. Morimoto and S. Serizawa. Strategy-proofness and efficiency with non-quasi-linear preferences: A characterization of minimum price walrasian rule. *Theoretical Economics*, 10(2): 445–487, 2015.

C. H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and system Sciences*, 48(3):498–532, 1994.

M. Quinzii. Core and competitive equilibria with indivisibilities. *International Journal of Game Theory*, 13(1):41–60, 1984.

A. E. Roth and A. Postlewaite. Weak versus strong domination in a market with indivisible goods. *Journal of Mathematical Economics*, 4(2):131–137, 1977.

A. Rubinstein. *Hardness of Approximation between P and NP*. Morgan & Claypool, 2019.

H. Saitoh and S. Serizawa. Vickrey allocation rule with income effect. *Economic Theory*, 35: 391–401, 2008.

H. E. Scarf. The computation of equilibrium prices: an exposition. *Handbook of mathematical economics*, 2:1007–1061, 1982.

L. Shapley and H. Scarf. On cores and indivisibility. *Journal of mathematical economics*, 1(1): 23–37, 1974.

L. S. Shapley and M. Shubik. The assignment game i: The core. *International Journal of game theory*, 1:111–130, 1971.

E. Sperner. *Neuer Beweis für die Invarianz der Dimensionszahl und des Gebietes*. 7. Vandenhoeck & Ruprecht, 1928.

W. Stromquist. How to cut a cake fairly. *The American Mathematical Monthly*, 87(8):640–644, 1980.

F. E. Su. Rental harmony: Sperner's lemma in fair division. *The American mathematical monthly*, 106(10):930–942, 1999.

L.-G. Svensson. Competitive equilibria with indivisible goods. *Zeitschrift für Nationalökonomie/Journal of Economics*, 44(4):373–386, 1984.