CONSTRUCTIVE RATIONALIZATION: LEARNING FROM COLLECTIVE REVEALED PREFERENCES

Anonymous authors
Paper under double-blind review

ABSTRACT

This paper presents a revealed preference approach for learning the collective consumption behavior. We introduce the Constructive Rationalization Method (CRM), which approximates the real market via a surrogate market of artificial consumers, called androids, with easy-to-compute demand functions. CRM uses observed aggregate demand and adds artificial consumers on the fly, while redistributing wealth under an empirical risk minimization principle. Unlike classical revealed preference approaches, CRM provides guarantees on the generalization risk for learning the aggregate demand function, while relying little on individual rationality. As an application, CRM can be used to provide reliable predictions for collective consumption behavior. Specifically, we show how to apply CRM to approximate allocations that are proportionally fair.

1 INTRODUCTION

Demand analysis lies at the heart of decision-making problems in operations research. In revenue management, we are interested in deciding the prices of products to offer customers with a view to maximizing expected revenues (Talluri & van Ryzin, 2004). In resource allocation, we design mechanisms that allocate scarce goods so that what each agent receives matches their preferences (Varian, 1974; Kelly, 1997). These downstream decisions rest on learning the demand and its sensitivity to control variables such as prices.

Broadly speaking, demand is only a revelation of economic primitives like individuals' preferences or utilities. For example, consider the canonical problem with a set of goods $[n] := \{1, \dots, n\}$, and a population of agents $[m] := \{1, \dots, m\}$ as consumers. If we assume the rationality of an agent i, then its demand $\mathbf{x}_i$ should be consistent with the behavior that optimizes its preference relations $\succeq_i$ over the goods. The aggregate demand $\mathbf{d}$ summarizes the demands of all agents.

While individual-level preference elicitation is appealing, it is sometimes impractical and ethically sensitive in large-scale economic systems. The number of individuals may be enormous, and collecting or analyzing individual preference data can raise privacy and regulatory concerns (EU, 2016). In many applications, the available data consist only of a finite collection of aggregate demand observations, with the identities of the contributing individuals anonymized or unobserved. This forms the key question addressed in this paper: **How to efficiently support decision-making processes based only on collective and observable consumption behavior?**

1.1 MOTIVATION AND OUR APPROACH

We focus on the demand of consumers with Walrasian competitive budgets. Assume that the goods are divisible, and each of them is associated with a price $p_j \in \mathbb{R}_+$ and a limited amount of supply. For each agent i, its utility $u_i \in \mathscr{U}$ is a function that encodes the preferences $\succeq_i$ and depends on the current market price $\mathbf{p}$ and its wealth function $w_i \in \mathscr{W}$, which as well depends on $\mathbf{p}$. The demand $\mathbf{x}_i \in \mathscr{X}$ is induced by the Utility Maximization Problem[1],

$$\mathbf{x}_i(\mathbf{p}, w_i(\mathbf{p})) = \arg\max_{\mathbf{x} \in \mathbb{R}^n_+} \; u_i(\mathbf{x}), \text{ s.t. } \langle \mathbf{p}, \mathbf{x} \rangle \leq w_i(\mathbf{p}). \tag{UMP}$$

[1] Strictly speaking, the existence of utility requires regularity on the preference relation $\succeq$.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Thus, the aggregate demand function of the market is

$$\mathbf{d}(\mathbf{p}) = \sum_{i=1}^{m} \mathbf{x}_i(\mathbf{p}, w_i(\mathbf{p})). \tag{1}$$

One natural approach to analyzing $\mathbf{d}$ is to combine the inferences for each agent by rationalizing individual demand. This task has been a major topic in revealed preference theory (Samuelson, 1948; Chambers & Echenique, 2016). Specifically, consider the case where for each agent i, K observed price-demand pairs $\mathbf{\Xi}_i = \{(\mathbf{p}_k, \mathbf{x}_{ik})\}_{k\in[K]}$ are provided. We hope to find such a $u_i : \mathbb{R}^n_+ \mapsto \mathbb{R}_+$ that the observed demand $\mathbf{x}_{ik}$ maximizes the corresponding UMP for all $k \in [K]$. This turns out to be possible only if the data $\mathbf{\Xi}_i$ satisfies the Generalized Axiom of Revealed Preference (GARP), i.e., there are no preference cycles. A test can be performed by the Afriat-Varian inequalities (Afriat, 1967; Varian, 1982):

$$\begin{aligned} u_{ik'} \le u_{ik} + \lambda_{ik}\langle \mathbf{p}_k, \mathbf{x}_{ik'} - \mathbf{x}_{ik}\rangle, \quad \forall k, k' \in [K], \\ (u_{ik}, \lambda_{ik}) \in \mathbb{R} \times \mathbb{R}_+, \quad \forall k \in [K]. \end{aligned} \tag{2}$$

The infeasibility of (2) denies the existence of a utility function; otherwise, any feasible solution implies that GARP holds, and the data is rationalized by the following concave function:

$$u_i(\mathbf{x}) = \min_{k\in[K]} \left\{u_{ik} + \lambda_{ik}\langle \mathbf{p}_k, \mathbf{x} - \mathbf{x}_{ik}\rangle\right\}, \tag{3}$$

which is Piecewise-Linear Concave (PLC) and monotone. To query the demand of a PLC agent like this, one needs to solve a linear programming problem. It is possible to use parametric models, but not surprisingly, preferences recovered therein are usually misspecified (Varian, 1990; Halevy et al., 2018).

Leaving aside the scalability and privacy concerns in modern economic systems, the value of individual rationality seems quite limited. For example, we are often interested in finding a Proportionally Fair Allocation (PFA) of goods to each of the m agents. Specifically, an allocation $(\mathbf{x}_1, ..., \mathbf{x}_m) \in \mathcal{X}$ is said to be proportionally fair if for all other feasible allocations $(\mathbf{y}_1, ..., \mathbf{y}_m) \in \mathcal{X}$, the following condition holds,

$$\sum_{i\in[m]} w_i \frac{u_i(\mathbf{y}_i) - u_i(\mathbf{x}_i)}{u_i(\mathbf{x}_i)} \le 0, \quad \mathcal{X} = \left\{ (\mathbf{x}_1, ..., \mathbf{x}_m) \in \mathbb{R}^{n\times m}_+ \mid \sum_{i\in[m]} \mathbf{x}_i \le \mathbf{1} \right\}, \tag{PFA}$$

where the weights w_i can be interpreted as the (constant) wealth of the agent i (Varian, 1974; Kelly, 1997; Bertsimas et al., 2011). PFA is a normative notion desirable for many applications, such as pricing and allocation in communication networks, spectrum management, online advertising, kidney transplantation, and electricity markets; see, e.g., Srikant (2004); Luo & Zhang (2008); Bertsimas et al. (2013); Azizan Ruhi et al. (2018); Bateni et al. (2022). Computationally, a PFA can be obtained by maximizing the weighted Nash Social Welfare function (Eisenberg & Gale, 1959) over the set of feasible allocations $\mathcal{X}$:

$$\prod_{i\in[m]} u_i(\mathbf{x}_i)^{w_i}, \tag{NSW}$$

which is a convex optimization problem. A dilemma arises if the Afriat–Varian rationalization (3) is used for this purpose. On the one hand, achieving an accurate utility or demand representation requires a sufficiently large number of observations. On the other hand, incorporating more observations increases the size of the resulting NSW maximization. Moreover, Afriat–Varian rationalization only interpolates the observed data – the required sample size may be unbounded. An alternative route is to first find a Walrasian Equilibrium (WE) price $\mathbf{p}$, at which the aggregate demand $\mathbf{d}$ clears the supply of the market:

$$\mathbb{R}^n_+ \ni \mathbf{p} \perp \big(\mathbf{1} - \mathbf{d}(\mathbf{p})\big) \in \mathbb{R}^n_+, \qquad \text{where } \mathbf{d} \text{ is defined in Eq. (1).} \tag{WE}$$

Then, the allocation is obtained by assigning each agent its demand at the equilibrium price. Unlike maximizing NSW as a centralized planner with full knowledge of $\{(u_i, w_i)\}_{i\in[m]}$, a WE can in principle be reached through a decentralized price-adjustment process that relies only on demand $\mathbf{d}(\mathbf{p})$. However, computing a WE is generally more difficult. For example, with constant wealth and PLC utilities here, finding an approximate WE is PPAD-hard (Chen & Teng, 2009). To keep this route tractable, we often restrict $\{u_i\}_{i\in[m]}$ to simple parametric utilities (Gao et al., 2021; Bateni et al., 2022; Jalota & Ye, 2025).

In view of the above obstacles, we consider utilizing the aggregate demand $\mathbf{d}$ holistically. Given access to historical prices and the corresponding aggregate demands $\{(\mathbf{p}_k, \mathbf{d}(\mathbf{p}_k))\}_{k=1}^K$, we present an approach that constructively rationalizes this collective information. Without identifying the *real* agents in the market, we build a surrogate market to imitate the collective behavior by iteratively adding *easy* artificial consumers, called *androids*, and redistributing the wealth among them to produce a proper approximation. While the real market might possess complex demand functions and nonlinear wealth, each android has an easy-to-compute demand function and linear wealth. To ensure approximation quality, we add androids on the fly until no improvement is detected.

1.2 Contributions

We propose the Constructive Rationalization Method (CRM, Algorithm 1) with a surrogate market to approximate the aggregate demand $\mathbf{d}$ of the real market under empirical risk minimization (ERM). Because CRM learns from only aggregate historical observations, it requires no individual-level historical data or knowledge of $\mathscr{U}$, $\mathscr{X}$, and $\mathscr{W}$. In this paper, we make the following contributions.

- **Learning from aggregate demand.** In comparison to the existing studies on rationality (Afriat, 1967; Varian, 1982) and learning of individual preferences (Beigman & Vohra, 2006; Balcan et al., 2014; Chambers et al., 2021), CRM is a revealed preference approach for *collective behavior* based on a novel convex optimization formulation. We show how to extend the original formulation to accommodate non-homothetic preferences, linear wealth, and extra supply. We conduct numerical experiments to validate CRM's predictive performance. At deployment, CRM can approximate PFA using one demand query per current agent, without eliciting utilities or requiring iterative individual communication.
- **Generalization and convergence guarantees.** We provide guarantees on the generalization risk (Theorem 4.1) that depend on the Rademacher complexity of the classes of androids used in the surrogate market. We also establish the convergence rate (Theorem 4.2); because finding a *good* android may be challenging, we provide guarantees for possibly inaccurate solutions.
- **Quantifying the hardness of the learning problem.** We show that the main obstacle lies in finding a *good* android that improves the current estimate. We demonstrate that even for simple utility classes like linear, Leontief, and CES utilities, this task is computationally hard (cf. Table 1). We pair these hardness results with globally certifiable formulations and practical heuristics.
- **When misspecification disappears, and when it does not.** Restricting CRM to a prescribed class of androids can introduce misspecification. We identify conditions under which the full universe of homothetic androids with linear wealth can eliminate this error (Theorem 4.3). Unfortunately, we show that only using simple android classes, including linear utilities, faces a sharp worst-case approximation limit (Theorem 4.4).

1.3 Related work

More broadly, our work is related to a long-standing theme in economics: what can be inferred from collective consumption behavior? Classical results show that aggregate demand inherits remarkably little structure from individual rationality (Sonnenschein, 1973b; Mantel, 1974; Debreu, 1974), while the decomposition of collective behavior into individual preferences becomes difficult when wealth effects are unknown (Browning & Chiappori, 1998; Chiappori & Ekeland, 1999a; Cherchye et al., 2011). The approach in this paper contributes an algorithmic tool for studying this fundamental topic. A detailed literature review is deferred to Section A.

2 Preliminaries

In this section, we formally define the problem setting and discuss a few elements for later discussion. The real market has $j = 1, ..., n$ divisible goods and $i = 1, ..., m$ agents. Each good j is associated with one unit of supply and a price p_j; the price lies in a compact convex set, $\mathbf{p} \in \mathcal{P} \subseteq \mathbb{R}^n_+$. Each agent i has a wealth function $w_i \in \mathscr{W} : \mathcal{P} \mapsto \mathbb{R}_+$, and a utility function $u_i \in \mathscr{U} : \mathbb{R}^n_+ \mapsto \mathbb{R}$. Each u_i is locally non-satiated and concave, and the wealth w_i is bounded. The total wealth in the market is $W(\mathbf{p}) = \sum_{i=1}^m w_i(\mathbf{p})$. As mentioned, the demand is $\mathbf{x}_i \in \mathscr{X}$ for each agent i, and the

aggregate demand function is given by:

$$\mathbf{d}(\mathbf{p}) = \sum_{i=1}^{m} \mathbf{x}_i(\mathbf{p}, w_i(\mathbf{p})), \quad \text{and} \quad \mathbf{x}_i(\mathbf{p}, w_i(\mathbf{p})) = \arg\max_{\mathbf{x}\in\mathbb{R}^n_+} u_i(\mathbf{x}) \ \text{ s.t. } \ \langle \mathbf{p}, \mathbf{x}\rangle \le w_i(\mathbf{p}).$$

A deterministic tie-breaking rule is applied whenever the demand is set-valued. Also, introduce the market excess demand function:

$$\mathbf{z}(\mathbf{p}) = \mathbf{d}(\mathbf{p}) - \mathbf{1}. \tag{4}$$

Now we define the inputs of our problem. We assume no a priori knowledge about the underlying agents; instead, the input of our problem is the sequence of price-demand pairs. Let the distribution of price be τ. Given i.i.d. samples $\mathbf{p}_k \sim \tau, k \in [K]$, the revealed preferences are presented by $\mathbf{\Xi} = \{(\mathbf{p}_k, \mathbf{d}_k)\}_{k=1}^K$ with $\mathbf{d}_k \equiv \mathbf{d}(\mathbf{p}_k)$. For simplicity, we assume $\mathbf{p}_k \in \Delta_n$ and $W(\mathbf{p}_k) = 1$ in the real market, which will be relaxed later.

To model the collective behavior of the real market, we introduce a surrogate market with T artificial consumers, whose wealth and preferences are determined by our construction. In the rest of this paper, we preserve the index i and the term "agent" for a consumer in the real market, while t and "android" for an artificial consumer.

2.1 Expenditure shares

For either an android or a real agent (suppressing the subscript for notational convenience), its demand function $\mathbf{x}(\mathbf{p}, w(\mathbf{p}))$ solves its own UMP. Instead of using the demand function explicitly, we define the expenditure share mapping as follows:

$$\boldsymbol{\gamma}(\mathbf{p}, w(\mathbf{p})) = \frac{\mathrm{diag}(\mathbf{p})\mathbf{x}(\mathbf{p}, w(\mathbf{p}))}{w(\mathbf{p})} \in \Delta_n, \tag{5}$$

indicating the share of total wealth spent on each good. The symbol Δ_n means the n-dimensional simplex: $\Delta_n = \{\mathbf{x} : \langle \mathbf{x}, \mathbf{1}\rangle = 1, \mathbf{x} \in \mathbb{R}^n_+\}$. One reason to do so is that the demand functions can be unbounded and non-Lipschitz near the boundary (e.g., zero prices) under Euclidean geometry (Cheung et al., 2018; Zhang et al., 2025). In the surrogate market, we limit an android's expenditure share from a prescribed universe $\mathscr{H} = \bigcup_{f\in[\bar{f}]} \mathscr{H}^f$, the union of finitely many component classes:

$$\mathscr{H}^f \subseteq \{\boldsymbol{\gamma} : \Delta_n \times \mathbb{R}_+ \to \Delta_n\}, \ \forall f \in [\bar{f}]. \tag{6}$$

We give some examples of $\mathscr{H}$ in Section B. In fact, we can simply write $\boldsymbol{\gamma}(\mathbf{p}) \equiv \boldsymbol{\gamma}(\mathbf{p}, w)$ independent of wealth $w(\mathbf{p})$ if the preferences are homothetic (Lemma B.1); otherwise, the expenditure share remains to be a joint function of price and wealth. In the real market, while the expenditure share $\boldsymbol{\gamma}_i$ of a real agent i is not accessible, we can still acquire the market expenditure share:

$$\mathbf{g} : \Delta_n \mapsto \Delta_n, \quad \mathbf{g}(\mathbf{p}) = \frac{\mathrm{diag}(\mathbf{p})\mathbf{d}(\mathbf{p})}{\sum_{i\in[m]} w_i(\mathbf{p})} = \mathrm{diag}(\mathbf{p})\mathbf{d}(\mathbf{p}) \in \Delta_n. \tag{7}$$

The fact that $\mathbf{g}(\mathbf{p}) \in \Delta_n$ is apparent since the budget of each agent is exhausted by non-satiation, viz, $\langle \mathbf{1}, \mathbf{g}(\mathbf{p})\rangle = W(\mathbf{p}) = 1$. Thereby, the revealed preferences can be represented as,

$$\mathbf{\Xi} = \{\mathbf{p}_k, \mathbf{g}_k\}_{k=1}^K, \quad \mathbf{g}_k = \mathrm{diag}(\mathbf{p}_k)\mathbf{d}_k, \ \forall k \in [K].$$

3 Constructive rationalization by the surrogate market

With a collection of androids drawn from the universe $\mathscr{H}$, we need a scheme to distribute the wealth among them. To begin with, assume that an android's preferences are homothetic and the wealth function is exogenous to the price. In this case, the aggregate demand lies in the convex hull of the demand functions of the androids, so does the expenditure share. We propose extensions to accommodate non-homothetic androids and endogenous wealth functions.

3.1 The wealth redistribution problem

Let the wealth in the surrogate market be distributed according to a nonnegative Borel probability measure ω supported on $\mathscr{H}$: $\omega \in \mathfrak{M}(\mathscr{H})$. We assume that $\mathscr{H}$ is nonempty and its evaluation maps

are Borel measurable. With homothetic androids, the expenditure share in the surrogate market can be written as a function $\mathbf{h} : \Delta_n \to \Delta_n$: $\mathbf{h}(\mathbf{p}) = \int_{\mathscr{H}} \boldsymbol{\gamma}(\mathbf{p})\, \mathrm{d}\omega$. Finitely supported measures yield $\mathbf{h} \in \mathrm{conv}(\mathscr{H})$ and suffice to approximate the optimal empirical loss. The goal is to minimize the mismatch between $\mathbf{g}(\cdot)$ for the real market and $\mathbf{h}(\cdot)$ for the surrogate market. Given the real market shares $\mathbf{g}_k$, we write the **W**ealth Redistribution in **P**rimal Form (WP) as follows,

$$L^* = \inf_{\mathbf{h}\in\mathrm{conv}(\mathscr{H})} \frac{1}{K} \sum_{k\in[K]} \| \mathbf{g}_k - \mathbf{h}(\mathbf{p}_k)\| = \inf_{\omega\in\mathfrak{M}(\mathscr{H})} \frac{1}{K} \sum_{k\in[K]} \left\| \mathbf{g}_k - \int_{\mathscr{H}} \boldsymbol{\gamma}(\mathbf{p}_k) \mathrm{d}\omega \right\|, \qquad \text{(WP)}$$

where $\|\cdot\|$ is any convex semi-norm on $\mathbb{R}^n$. In view of the second equation, it is an infinite-dimensional convex optimization problem over the set of probability measures $\mathfrak{M}(\mathscr{H})$. Here, the mismatch between two markets is measured through a scaled local metric $\|\mathrm{diag}(\mathbf{p}_k)(\cdot)\|$.

The dual problem, **W**ealth Redistribution in **D**ual Form, denoted WD, is the following (strong duality holds by Lemma C.1):

$$L^* = \max_{\mathbf{U},\mu} \frac{1}{K} \left(\sum_{k\in[K]} \langle \mathbf{U}\mathbf{e}_k, \mathbf{g}_k \rangle - \mu \right), \quad \text{s.t. } (\mathbf{U}, \mu) \in \mathcal{F}(\mathscr{H}), \qquad \text{(WD)}$$

where the dual feasible set, associated with the universe $\mathscr{H}$, is given by,

$$\mathcal{F}(\mathscr{H}) = \left\{ (\mathbf{U}, \mu) \in \mathbb{R}^{Kn} \times \mathbb{R} : \sum_{k\in[K]} \langle \mathbf{U}\mathbf{e}_k, \boldsymbol{\gamma}(\mathbf{p}_k) \rangle \le \mu,\ \forall \boldsymbol{\gamma} \in \mathscr{H};\ \|\mathbf{U}\mathbf{e}_k\|_* \le 1,\ \forall k \in [K] \right\}. \quad (8)$$

Here, $\|\cdot\|_*$ denotes the dual norm, and $\mathbf{e}_k$ is the k-th canonical basis of $\mathbb{R}^K$.

3.2 The conceptual algorithm

We are now ready to present the conceptual method (Algorithm 1). The algorithm starts with T_0 androids and iteratively adds new androids to the collection and update the wealth distribution. With T androids, the measure to distribute the wealth becomes a T-dimensional vector $\mathbf{w} \in \Delta_T$, and so $\mathbf{h}(\mathbf{p}) = \sum_{t=1}^{T} w_t \boldsymbol{\gamma}_t(\mathbf{p}, w_t)$.

We use the subscript $t \in [T]$ to denote related quantities of the t-th android. Based on the discussion in Section 2.1, the t-th android will be denoted by its expenditure share $\boldsymbol{\gamma}_t(\mathbf{p}) \in \mathscr{H}$. When wealth redistribution over the current T androids is solved exactly, the risk equals the value of the following finite primal-dual pair:

$$\begin{aligned} &\text{(WP}_T\text{)} && \text{(WD}_T\text{)} \\ L_T^* = \min_{\mathbf{w}\in\Delta_T} \tfrac{1}{K} \textstyle\sum_{k\in[K]} \|\mathbf{g}_k - \mathbf{h}(\mathbf{p}_k)\| &= \max_{\substack{\|\mathbf{U}\mathbf{e}_k\|_* \le 1 \\ \boldsymbol{\mu}\in\mathbb{R}^n}} \tfrac{1}{K} \left(\textstyle\sum_{k\in[K]} \langle \mathbf{U}\mathbf{e}_k, \mathbf{g}_k \rangle - \mu \right) \\ &&\text{s.t. } \textstyle\sum_{k\in[K]} \langle \mathbf{U}\mathbf{e}_k, \boldsymbol{\gamma}_t(\mathbf{p}_k) \rangle - \mu \le 0,\ \forall t \in [T]. \end{aligned}$$

When updating $\{w_t(\cdot)\}_{t=1}^T$, one can think of some lump-sum payments being transferred among the androids. To find the new android, we rely on the following oracle called SEP, provided a linear operator $\mathbf{V} \in \mathbb{R}^{Kn}$,

$$\pi_{T+1} = \sup_{\boldsymbol{\gamma}\in\mathscr{H}} \sum_{k=1}^{K} \langle \mathbf{V}\mathbf{e}_k, \boldsymbol{\gamma}(\mathbf{p}_k) \rangle. \qquad \text{(SEP)}$$

Algorithm 1: CRM, Constructive Rationalization Method (by the Surrogate Market)

Input: Dataset $\boldsymbol{\Xi} = \{\mathbf{p}_k, \mathbf{g}_k\}_{k=1}^K$; initial android collection $\{\boldsymbol{\gamma}_t\}_{t=1}^{T_0}$, $T_0 \ge 1$.

for $T = T_0, T_0 + 1, \ldots$ **do**

  Update wealth distribution $\{w_t(\cdot)\}_{t=1}^T$;

  Find the new android ($\boldsymbol{\gamma}_{T+1}$) by solving (SEP);

  Add android ($\boldsymbol{\gamma}_{T+1}$) to the collection; `// Terminate if no/small improvement.`

end

Output: Surrogate market $\{(\boldsymbol{\gamma}_t(\cdot), w_t(\cdot))\}_{t=1}^T$.

Note CRM is a conceptual method because it does not specify how to update the wealth distribution $\{w_t(\cdot)\}_{t=1}^T$ and how the oracle is solved (e.g., choice of $\mathbf{V}$ and $\mathscr{H}$). For example, we can at least give two concrete implementations.

Implementation 1: Cutting Plane Method (CPM). To implement CRM as a CPM, we solve (WP_T) or (WD_T) to obtain dual solutions $(\mathbf{U}_T, \mu_T)$. Then, we may define $\mathbf{V} \leftarrow \mathbf{U}_T$ and use (SEP) to find the new android. If so, the oracle is equivalent to a Separation Problem (thus qualifying the name SEP). If $\pi_{T+1} > \mu_T$, we add the corresponding android with spending share function $\boldsymbol{\gamma}_{T+1}$ to the market; otherwise, we can stop the algorithm. It is also possible to choose the dual solutions as certain geometric centers; see, e.g., Grötschel et al. (1993); Vaidya (1996).

Implementation 2: Frank-Wolfe (FW) method . It is also possible to implement CRM as a FW method (Frank & Wolfe, 1956; Jaggi, 2013; Lacoste-Julien & Jaggi, 2015) that uses the oracle as a linear optimization oracle over the space $\mathrm{conv}(\mathscr{H})$. In this paper, we use this specific implementation and derive the convergence rate of CRM. To maintain the flow, we leave details to Section E.

Furthermore, CRM can be seen as a boosting-type algorithm (Bartlett et al., 1998; Demiriz et al., 2002) to solve the empirical risk minimization. Specifically, for each android t, its expenditure share $\boldsymbol{\gamma}_t(\mathbf{p})$ can be seen as a (weak) learner of $\mathbf{g}(\mathbf{p})$. It follows that $\mathscr{H}$ is the *hypothesis class* of the base learners, the final predictor satisfies $\mathbf{h} \in \mathrm{conv}(\mathscr{H})$. For each sample k, the loss is computed by:

$$\ell(\mathbf{h}(\mathbf{p}_k), \mathbf{g}_k) := \|\mathbf{h}(\mathbf{p}_k) - \mathbf{g}_k\|. \tag{10}$$

The empirical risk of the current predictor after T iterations is

$$L(\mathbf{h}) := \tfrac{1}{K} \sum_{k=1}^{K} \ell(\mathbf{h}(\mathbf{p}_k), \mathbf{g}_k) = L_T. \tag{11}$$

The above viewpoint helps to derive a generalization bound at unseen prices. The equivalence to a boosting algorithm also clears the apprehension that a large number of androids may deteriorate the generalization risk. It is also worth noting that non-homothetic androids can be handled using the following strategy.

Remark 1. *Directly using non-homothetic androids makes WP_T nonconvex, because $\mathbf{h} = \sum_{t\in[T]} \boldsymbol{\gamma}_t(\mathbf{p}_k, w_t) w_t$ is nonlinear in w_t, and thus only weak duality holds. To simplify, we can adopt a homothetic block H and a non-homothetic block N. Each non-homothetic android $t \in N$ is added with a prefixed budget $w_0 \in (0, 1]$ (thus $|N| \le \lfloor \frac{1}{w_0} \rfloor$). Namely, WP_T becomes:*

$$L_T^* = \min_{\mathbf{w}\in\Delta_T, w_t = w_0,\ \forall t\in N} \tfrac{1}{K} \sum_{k=1}^{K} \left\| \sum_{t\in H} w_t\, \boldsymbol{\gamma}_t(\mathbf{p}_k) + \sum_{t\in N} w_t\, \boldsymbol{\gamma}_t(\mathbf{p}_k, w_0) - \mathbf{g}_k \right\|.$$

The same strategy can be used to solve the separation problem as before.

3.3 Extensions to linear wealth and extra supply

We now consider an extension where androids adopt linear wealth functions and a single supplier provides an extra supply of goods $\mathbf{a}$. Then, redistribute the enlarged supply among the androids as endowments $\mathbf{b}_t$, so that $w_t(\mathbf{p}) = \langle \mathbf{p}, \mathbf{b}_t \rangle$, and $\sum_{t\in[T]} \mathbf{b}_t = \mathbf{1} + \mathbf{a}, \mathbf{b}_t \ge \mathbf{0}$. We consider the following class of predictors,

$$\mathbf{h}(\mathbf{p}) := \sum_{t\in[T]} \langle \mathbf{p}, \mathbf{b}_t \rangle \boldsymbol{\gamma}_t(\mathbf{p}) - \mathrm{diag}(\mathbf{p})\mathbf{a}; \tag{12}$$

We denote this predictor class by $\mathscr{Q}$. Unlike the constant-wealth predictors, a predictor $\mathbf{h} \in \mathscr{Q}$ need not lie in $\mathrm{conv}(\mathscr{H})$. For T androids, we optimize the extra supply jointly with the android endowments[2]:

$$L_{+,T}^* = \underset{\substack{\mathbf{a},\mathbf{b}_1,\dots,\mathbf{b}_T \\ \text{s.t. } \mathbf{a}\ge\mathbf{0},\quad \mathbf{b}_t\ge\mathbf{0},\quad t\in[T], \\ \sum_{t\in[T]}\mathbf{b}_t = \mathbf{1}+\mathbf{a},}}{\min} \tfrac{1}{K}\sum_{k\in[K]} \|\mathbf{g}_k - \mathbf{h}(\mathbf{p}_k)\| \quad (\mathrm{WP}_{+,T}) = \underset{\substack{\mathbf{U},\boldsymbol{\mu} \\ \text{s.t. } \|\mathbf{U}\mathbf{e}_k\|_* \le 1,\quad k\in[K],\quad \boldsymbol{\mu}\in\mathbb{R}^n, \\ \sum_{k\in[K]}\langle \mathbf{U}\mathbf{e}_k, \boldsymbol{\gamma}_t(\mathbf{p}_k)\rangle \mathbf{p}_k \le \boldsymbol{\mu},\quad t\in[T], \\ \boldsymbol{\mu} \le \sum_{k\in[K]} \mathrm{diag}(\mathbf{p}_k)\mathbf{U}\mathbf{e}_k.}}{\max} \tfrac{1}{K}\left(\sum_{k\in[K]} \langle \mathbf{U}\mathbf{e}_k, \mathbf{g}_k\rangle - \langle \boldsymbol{\mu}, \mathbf{1}\rangle\right) \quad (\mathrm{WD}_{+,T})$$

[2]The symbol $+$ marks that the problem is an extension to the original WP_T.

Both remain convex optimization problems. Minimizing over $\mathbf{a} \geq \mathbf{0}$ introduces the last dual constraint, which is independent of the android collection. Setting $\mathbf{a} = \mathbf{0}$, $\mathbf{b}_t = w_t\mathbf{1}$ recovers WP_T, so clearly, $L^*_{+,T} \leq L^*_T$. Similarly, for a prescribed linear operator $\mathbf{V} \in \mathbb{R}^{n\times K}$, a new android is found by solving

$$\pi_{+,T+1} = \max_{j\in[n]} \left\{ \max_{\boldsymbol{\gamma}\in\mathscr{H}} \sum_{k=1}^{K} \langle p_{k,j}\mathbf{V}\mathbf{e}_k, \boldsymbol{\gamma}(\mathbf{p}_k)\rangle - \mu_{T,j} \right\}, \tag{14}$$

which is equivalent to solve SEP for each $j \in [n]$.

4 Performance Guarantees

In this section, we discuss the performance guarantees of CRM. For comprehensiveness, we derive the results for $\mathrm{WP}_{+,T}$ so that the original WP_T may be taken as a special case. Throughout, we assume that the number of classes forming universe $\mathscr{H} = \bigcup_{f=1}^{\bar{f}} \mathscr{H}^f$ is finite, $\bar{f} < \infty$, and that the extra supply satisfies $\mathbf{0} \leq \mathbf{a} \leq D_{\mathbf{a}}\mathbf{1}$ for some fixed $0 \leq D_{\mathbf{a}} < \infty$. For simplicity, we specify the norm used in CRM to be $\|\cdot\| \equiv \|\cdot\|_\infty$.

4.1 Generalization Risk and Convergence

We first state the following generalization bound.

Theorem 4.1 (Generalization bound)**.** *For any $\delta \in (0,1)$, with probability at least $1-\delta$ over the i.i.d. sampled prices $\mathbf{p}_1, \ldots, \mathbf{p}_K \sim \tau$, the following bound holds for every $\mathbf{h} \in \mathscr{Q}$:*

$$\begin{aligned}\mathbb{E}_\tau[\ell(\mathbf{h}(\mathbf{p}), \mathbf{g}(\mathbf{p}))] \leq L(\mathbf{h}) + 2D_{\mathbf{a}}\sqrt{\frac{n}{K}} \\ + (1+D_{\mathbf{a}})\left(4n\sqrt{\frac{\log \bar{f}}{K}} + 2\sqrt{2}n \max_{f\in[\bar{f}]} \hat{\mathcal{R}}_K(\mathscr{H}^f) + 6\sqrt{\frac{\ln(2/\delta)}{2K}}\right),\end{aligned} \tag{15}$$

where $\hat{\mathcal{R}}_K(\mathscr{H}^f)$ is the empirical Rademacher complexity of $\mathscr{H}^f$.

We leave the proof to Section D. Constant-wealth predictors from WP_T are obtained by taking $D_{\mathbf{a}} = 0$ and $\mathbf{b}_t = w_t\mathbf{1}$. This specialization removes the explicit prefactor n and yields the sharper bound in Theorem D.1. Although $\mathscr{Q}$ is not itself a convex hull, its decomposition into n convex-hull components preserves the key benefit of boosting: the generalization gap is independent of the number of androids. For the parametric utility classes in Example 1, under Assumption D.1, we could further show that the risk is of order $\widetilde{\mathcal{O}}(K^{-1/2})$.

We implement CRM as a Frank-Wolfe method. For the extension $\mathrm{WP}_{+,T}$, each linear minimization step reduces to n calls to SEP and a coordinatewise comparison of the supplier's two endpoint choices. Because the empirical loss is nonsmooth, we analyze a smoothed objective following Nesterov (2005). The formal construction and proof are given in Section E. The empirical risk of the iterate after T iterations can be decomposed into three parts,

$$L_{+,T} \leq \mathcal{O}(T^{-1/2}) + E_T + L^*_+. \tag{16}$$

The first term indicates the convergence rate of the Frank-Wolfe method and the smoothing error. The second term E_T is the weighted average error from solving SEP. The last term L^*_+ is the misspecification error, i.e., the smallest empirical risk achievable over $\mathscr{Q}$.

Theorem 4.2 (Convergence of CRM)**.** *Implement CRM over $\mathscr{Q}$ as in Section E.3. After $T \geq 1$ iterations, the output risk satisfies*

$$0 \leq L_{+,T} - L^*_+ \leq E_T + \frac{2\sqrt{2}(1+2D_{\mathbf{a}})}{\sqrt{T+2}}, \tag{17}$$

If the unrestricted formulation $\mathrm{WP}_{+,T}$ has an optimizer $\mathbf{a}^*$, then every $D_{\mathbf{a}} \geq \|\mathbf{a}^*\|_\infty$ retains an unrestricted optimizer, so L^*_+ is also the unrestricted optimum. Next, we discuss the two remaining terms in (16): the error from the SEP calls, E_T, and the misspecification error, L^*_+.

4.2 Hardness and solvability of the separation problem

We consider E_T in the convergence analysis because, unlike standard convex optimization over a "simple" set where a separating hyperplane is easy to find (e.g., from the subdifferential), CRM (implemented either by a CPM or a Frank-Wolfe method) requires solving the generally nonconvex separation problem SEP. We show that solving it is generally computationally hard, even for simple cases such as CES and Leontief (see Example 1). The efficiency of solving SEP is summarized in Table 1; details are given in Section H.

Table 1: Complexity and solution methods by oracle class. The FPTAS here means for fixed K.

Class	$K = 1$	K (part of input)	Method for $K \geq 2$
Cobb-Douglas	Polynomial-time	Polynomial-time	-
Linear		Strongly NP-hard (Theorem H.1)	MIP
CES/Leontief		NP-hard (Theorem H.2)	FPTAS (global, Theorem H.3) Spatial B&B (global, Section H.4) FPTAS (KKT, Remark 7)

For linear oracles, we can solve SEP by enumerating n^K choice patterns, which takes polynomial time for fixed K; however, by a log-price formulation we show it is strongly NP-hard when K is part of the input (Theorem H.1). It is also possible to use a mixed-integer programming (MIP) solver with the margin test and exclusion cuts; see Section H.1.

For CES oracles, we provide a global approximation guarantee for fixed K, but its running time scales exponentially with K (Theorem H.3). We can also solve a nonlinear formulation by spatial branch-and-bound method (Section H.4), which provides a certificate showing a feasible value and a global upper bound, and these errors determine E_T; however, we later see that tightening the upper bound can be costly (cf. Section I.2.3). To be practical, we also discuss a local scheme obtaining an ϵ-approximate KKT point (cf. Remark 7).

At a high level, the reduction is in the same spirit of axioms of revealed preference (Varian, 1982) that a rational consumer must admit no preference cycles, and the hardness result to rationalize household choices (Nobibon et al., 2016). To this end, the hardness of solving SEP globally is likely to be true for more complex classes.

4.3 Misspecification error

What remains in the risk estimates is the misspecification error L_+^*. How should the universe $\mathscr{H}$ and the wealth functions be chosen $\mathscr{W}$ when running CRM? The positive news is if we use linear wealth functions and $\mathscr{H}$ that includes *all* homothetic classes, the misspecification error could be eliminated if we use $\mathrm{WP}_{+,T}$. The formal statement and proof are given in Section F.

Theorem 4.3 (Informal)**.** *For a real market with linear wealth and smooth, degree-zero homogeneous excess demand satisfying Walras's law, homothetic androids with linear wealth achieve zero misspecification error on compact interior price domains if $WP_{+,T}$ is used.*

The proof of Theorem 4.3 provides a finite correction $\mathbf{a} = \varsigma\mathbf{1}$. The common scale $1 + \varsigma$ is chosen sufficiently large relative to the curvature bound in Assumption F.2 and the compact interior price domain. Hence, taking $D_{\mathbf{a}} \geq \varsigma$ and letting $\mathscr{H}$ contain all homothetic expenditure-share functions gives $L_+^* = 0$ in Theorem 4.2, and therefore

$$L_{+,T} \leq E_T + \frac{2\sqrt{2}(1 + 2D_{\mathbf{a}})}{\sqrt{T+2}}. \tag{18}$$

Thus the representation and convergence results concern the same linear-wealth and extra-supply formulation; the representation result does not by itself provide an efficient SEP oracle for the full homothetic universe. Simple parametric classes such as those in Example 1 can nevertheless face limits in uniform approximation over the entire open price simplex.

Theorem 4.4 (Informal)**.** *There exists a two-good, one-agent homothetic market with share* $\mathbf{g}$ *for which mixtures of androids with uniformly bounded share variation and linear wealth satisfy*

$$\inf_{\mathbf{h}} \sup_{\mathbf{p}\in\mathrm{int}\Delta_2} \|\mathbf{h}(\mathbf{p}) - \mathbf{g}(\mathbf{p})\|_\infty \geq \frac{\tanh(1)}{2} \approx 0.3808. \tag{19}$$

The bound is sharp and attainable by using only Cobb-Douglas consumers.

The formal statement and proof in Section G cover a broader class of polynomial wealth functions. The result above is a structural lower bound for uniform approximation over the open simplex. It does not imply a positive empirical loss on a fixed sample or a positive expected loss under a given price distribution, because the poorly approximated price may depend on the surrogate and may receive negligible probability under that distribution.

5 Numerical Experiments

After ablating the FW variants, problem size, and separation solver, we evaluate CRM in two applications using stochastic fully corrective Frank–Wolfe (`SFCFW`) and its Arrow–Debreu (AD) extension. Complete experimental configurations and results are provided in Section I.

5.1 Application I: Prediction under different wealth functions

Our first application evaluates CRM's predictive performance across market types and wealth specifications. The results are shown below.

Table 2: CRM performance. Mean ± std over five seeds; L_∞ in 10^{-3}.

		CES			GES		
Real wealth	Method	T	Train L_∞	Test L_∞	T	Train L_∞	Test L_∞
Constant	`SFCFW`	80±17	0.41±0.45	**0.48±0.47**	500	10.99±0.54	23.79±1.76
	`SFCFW(AD)`	62±34	**0.29±0.18**	0.53±0.31	126±23	**5.73±0.47**	**20.18±1.09**
First-order	`SFCFW`	81±13	20.20±2.80	22.70±3.20	500	10.92±1.08	23.14±2.01
	`SFCFW(AD)`	96±19	**0.24±0.13**	**0.61±0.33**	87±2	**6.46±0.16**	**20.25±0.69**
Second-order	`SFCFW`	59±4	52.30±2.70	56.80±3.80	500	13.82±1.31	28.38±1.14
	`SFCFW(AD)`	225±23	**22.50±3.10**	**28.30±3.80**	117±5	**7.07±0.81**	**22.77±1.01**

For constant wealth, SFCFW has lower mean test risk on CES, whereas SFCFW(AD) has lower mean test risk on GES; for first- and second-order wealth, SFCFW(AD) is lower on both markets.

5.2 Application II: Computing fair allocations

Our second application targets fair allocation when individual utilities are unavailable. CRM learns a surrogate market from historical aggregate demand; at deployment, its market-clearing price requires only one demand query per participant, after which we rescale the reports to satisfy supply. In CES and non-homothetic GES experiments with 1,000–5,000 agents and 10–100 goods, the resulting allocations have mean weighted Nash-welfare loss below 0.2%, using about 135–242 active androids. Under bounded elasticity, we also establish a bound on this welfare loss (Theorem I.1).

6 Conclusion

We study the Constructive Rationalization Method (CRM), which assembles a surrogate market from simple androids to learn collective revealed preferences. Because the surrogate is itself a market, it is directly amenable to downstream analytics. One can implement the equilibrium price of it to the real market for an allocation that is approximately fair. Several questions remain open. As we show that there is a limit to using simple Androids like the linear or CES classes, it is not clear whether we can find a simple parametric class sufficient to approximate a general market. On the computational side, it is also our future work to strengthen the current implementation, which requires better separation oracles and stochastic approximation schemes.

AI USE STATEMENT

Generative AI was used to assist with converting the existing manuscript to the ICLR 2027 LaTeX format, adapting the LaTeX preamble, checking compilation and cross-references, and help with the grammar and typos.

REPRODUCIBILITY STATEMENT

The main text states the learning problem, assumptions, algorithms, and theoretical guarantees. The appendices derive the demand functions and expenditure shares (Section B) and the primal-dual formulation (Section C); prove the generalization guarantees (Section D) and the convergence of CRM with inexact separation oracles (Section E); and provide the separation-oracle algorithms and complexity analysis (Section H) and the misspecification lower bound (Section G). The main text summarizes the numerical findings (Section 5), while Section I reports the market-generation procedures, implementation and solver settings, evaluation protocol, and full multi-seed results.

REFERENCES

S. N. Afriat. The construction of utility functions from expenditure data. *International Economic Review*, 8(1):67–77, 1967.

Kenneth J. Arrow and Leonid Hurwicz. On the stability of the competitive equilibrium, I. *Econometrica*, 26(4):522–552, 1958.

Navid Azizan Ruhi, Krishnamurthy Dvijotham, Niangjun Chen, and Adam Wierman. Opportunities for price manipulation by aggregators in electricity markets. *IEEE Transactions on Smart Grid*, 9(6):5687–5698, 2018.

Maria-Florina Balcan, Amit Daniely, Ruta Mehta, Ruth Urner, and Vijay V. Vazirani. Learning economic parameters from revealed preferences, 2014. URL http://arxiv.org/abs/1407.7937.

Peter Bartlett, Yoav Freund, Wee Sun Lee, and Robert E. Schapire. Boosting the margin: a new explanation for the effectiveness of voting methods. *The Annals of Statistics*, 26(5):1651–1686, 1998.

Peter L. Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: risk bounds and structural results. *J. Mach. Learn. Res.*, 3(11):463–482, 2003.

Pathikrit Basu and Federico Echenique. On the falsifiability and learnability of decision theories. *Theoretical Economics*, 15(4):1279–1305, 2020.

MohammadHossein Bateni, Yiwei Chen, Dragos Florin Ciocan, and Vahab Mirrokni. Fair resource allocation in a volatile marketplace. *Operations Research*, 70(1):288–308, 2022.

Xiaohui Bei, Wei Chen, Jugal Garg, Martin Hoefer, and Xiaoming Sun. Learning market parameters using aggregate demand queries. *Proceedings of the AAAI Conference on Artificial Intelligence*, 30(1), 2016.

Eyal Beigman and Rakesh Vohra. Learning from revealed preference. In *Proceedings of the 7th ACM conference on Electronic commerce*, EC '06, pp. 36–42, New York, NY, USA, 2006. Association for Computing Machinery.

Dimitris Bertsimas, Vivek F. Farias, and Nikolaos Trichakis. The price of fairness. *Operations Research*, 59(1):17–31, 2011.

Dimitris Bertsimas, Vivek F. Farias, and Nikolaos Trichakis. Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research*, 61(1):73–87, 2013.

Wei Bian, Xiaojun Chen, and Yinyu Ye. Complexity analysis of interior point algorithms for non-lipschitz and nonconvex minimization. *Mathematical Programming*, 149(1):301–327, 2015.

Richard Blundell, Martin Browning, and Ian Crawford. Best nonparametric bounds on demand responses. *Econometrica*, 76(6):1227–1262, 2008.

Richard Blundell, Martin Browning, Laurens Cherchye, Ian Crawford, Bram De Rock, and Frederic Vermeulen. Sharp for SARP: nonparametric bounds on counterfactual demands. *American Economic Journal: Microeconomics*, 7(1):43–60, 2015.

Richard W. Blundell, Martin Browning, and Ian A. Crawford. Nonparametric engel curves and revealed preference. *Econometrica*, 71(1):205–240, 2003.

Donald J. Brown, Peter M. Demarzo, and B. Curtis Eaves. Computing equilibria when asset markets are incomplete. *Econometrica*, 64(1):1–27, 1996.

M. Browning and P. A. Chiappori. Efficient intra-household allocations: a general characterization and empirical tests. *Econometrica*, 66(6):1241–1278, 1998.

Christopher P. Chambers and Federico Echenique. *Revealed Preference Theory*. Econometric society monographs. Cambridge University Press, United Kingdom ; New York, 2016.

Christopher P. Chambers, Federico Echenique, and Nicolas S. Lambert. Recovering preferences from finite data. *Econometrica*, 89(4):1633–1664, 2021.

A. Charnes and W. W. Cooper. Programming with linear fractional functionals. *Naval Research Logistics Quarterly*, 9(3-4):181–186, 1962.

Xi Chen and Shang-Hua Teng. Spending is not easier than trading: on the computational equivalence of Fisher and Arrow-Debreu equilibria. In Yingfei Dong, Ding-Zhu Du, and Oscar Ibarra (eds.), *Algorithms and Computation*, pp. 647–656, Berlin, Heidelberg, 2009. Springer.

Xi Chen, Dimitris Paparas, and Mihalis Yannakakis. The complexity of non-monotone markets. *J. ACM*, 64(3):20:1–20:56, 2017.

Laurens Cherchye, Bram De Rock, Jeroen Sabbe, and Frederic Vermeulen. Nonparametric tests of collectively rational consumption behavior: An integer programming procedure. *Journal of Econometrics*, 147(2):258–265, 2008.

Laurens Cherchye, Bram De Rock, and Frederic Vermeulen. The revealed preference approach to collective consumption behaviour: testing and sharing rule recovery. *The Review of Economic Studies*, 78(1):176–198, 2011.

Laurens Cherchye, Thomas Demuynck, Bram De Rock, and Joshua Lanier. Are consumers (approximately) rational? Shifting the burden of proof. *The Review of Economics and Statistics*, 107(6): 1652–1666, 2025.

Yun Kuen Cheung, Richard Cole, and Yixin Tao. Dynamics of distributed updating in Fisher markets. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, EC '18, pp. 351–368, New York, NY, USA, 2018. Association for Computing Machinery.

Yun Kuen Cheung, Richard Cole, and Nikhil R. Devanur. Tatonnement beyond gross substitutes? Gradient descent to the rescue. *Games and Economic Behavior*, 123:295–326, 2020.

P. A. Chiappori and I. Ekeland. Aggregation and market demand: an exterior differential calculus viewpoint. *Econometrica*, 67(6):1435–1457, 1999a.

P. A. Chiappori and I. Ekeland. Disaggregation of excess demand functions in incomplete markets. *Journal of Mathematical Economics*, 31(1):111–129, 1999b.

P. A. Chiappori and Jean-Charles Rochet. Revealed preferences and differentiable demand. *Econometrica*, 55(3):687–691, 1987.

Bruno Codenotti and Kasturi Varadarajan. Computation of market equilibria by convex programming. In Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani (eds.), *Algorithmic Game Theory*, pp. 135–158. Cambridge University Press, 2007.

Thomas M. Cover. Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition. *IEEE Transactions on Electronic Computers*, EC-14(3):326–334, 1965.

Gerard Debreu. Excess demand functions. *Journal of Mathematical Economics*, 1(1):15–21, 1974.

Ayhan Demiriz, Kristin P. Bennett, and John Shawe-Taylor. Linear programming boosting via column generation. *Machine Learning*, 46(1):225–254, 2002.

Daniele Depetrini and Marco Locatelli. Approximation of linear fractional-multiplicative problems. *Mathematical Programming*, 128(1):437–443, 2011.

N.R. Devanur, C.H. Papadimitriou, A. Saberi, and V.V. Vazirani. Market equilibrium via a primal-dual-type algorithm. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, pp. 389–395, 2002.

John Duchi. Statistics 231 / CS229T: Statistical Learning Theory, 2017. URL https://web.stanford.edu/class/cs229t/2017/homework.html.

Federico Echenique, Sangmok Lee, and Matthew Shum. The money pump as a measure of revealed preference violations. *Journal of Political Economy*, 119(6):1201–1223, 2011.

Edmund Eisenberg and David Gale. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics*, 30(1):165–168, 1959.

EU. Regulation (EU) 2016/679 of the European Parliament and of the Council. *Official Journal of the European Union*, 2016.

G. B. Folland. *Real Analysis: Modern Techniques and Their Applications*. Pure and Applied Mathematics. Wiley, New York, 2nd ed edition, 1999.

Marguerite Frank and Philip Wolfe. An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3(1-2):95–110, 1956.

Roland W. Freund and Florian Jarre. Solving the sum-of-ratios problem by an interior-point method. *Journal of Global Optimization*, 19(1):83–102, 2001.

Yuan Gao, Alex Peysakhovich, and Christian Kroer. Online market equilibrium with application to fair division. In *Advances in Neural Information Processing Systems*, volume 34, pp. 27305–27318, 2021.

M. R. Garey, D. S. Johnson, and L. Stockmeyer. Some simplified NP-complete graph problems. *Theoretical Computer Science*, 1(3):237–267, 1976. doi: 10.1016/0304-3975(76)90059-1.

Jugal Garg and László A. Végh. A strongly polynomial algorithm for linear exchange markets. *Operations Research*, 71(2):487–505, 2023.

Jugal Garg, Yixin Tao, and László A. Végh. Approximating competitive equilibrium by Nash welfare. *Games and Economic Behavior*, 158:142–166, 2026.

Dongdong Ge, Qi Huangfu, Zizhuo Wang, Jian Wu, and Yinyu Ye. Cardinal optimizer (COPT) user guide, 2022. URL http://arxiv.org/abs/2208.14314.

J. D. Geanakoplos and H. M. Polemarchakis. On the disaggregation of excess demand functions. *Econometrica*, 48(2):315–331, 1980.

Denizalp Goktas, Jiayi Zhao, and Amy Greenwald. Tâtonnement in homothetic Fisher markets. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*, pp. 760–781. ACM, 2023.

Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2 of *Algorithms and Combinatorics*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1993. URL http://link.springer.com/10.1007/978-3-642-78240-4.

Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual: Nonlinear constraints, 2026. URL https://docs.gurobi.com/projects/optimizer/en/current/features/nonlinear.html. Accessed September 22, 2026.

Yoram Halevy, Dotan Persitz, and Lanny Zrill. Parametric recoverability of preferences. *Journal of Political Economy*, 126(4):1558–1593, 2018.

Martin Jaggi. Revisiting frank-wolfe: Projection-free sparse convex optimization. In *International Conference on Machine Learning (ICML)*, pp. 427–435, 2013.

Kamal Jain. A polynomial time algorithm for computing an Arrow–Debreu market equilibrium for linear utilities. *SIAM Journal on Computing*, 37(1):303–318, 2007.

Kamal Jain and Vijay V. Vazirani. Eisenberg–Gale markets: Algorithms and game-theoretic properties. *Games and Economic Behavior*, 70(1):84–106, 2010.

Devansh Jalota and Yinyu Ye. Stochastic online fisher markets: static pricing limits and adaptive enhancements. *Operations Research*, 73(2):798–818, 2025.

Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller, James W. Thatcher, and Jean D. Bohlinger (eds.), *Complexity of Computer Computations*, pp. 85–103. Springer, Boston, MA, 1972.

Frank Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8(1):33–37, 1997.

Simon Lacoste-Julien and Martin Jaggi. On the global linear convergence of Frank-Wolfe optimization variants. *Advances in Neural Information Processing Systems*, 28, 2015.

Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces: Isoperimetry and Processes*. Springer Science & Business Media, 1991.

Zhi-Quan Luo and Shuzhong Zhang. Dynamic spectrum management: Complexity and duality. *IEEE Journal of Selected Topics in Signal Processing*, 2(1):57–73, 2008.

Rolf R Mantel. On the characterization of aggregate excess demand. *Journal of Economic Theory*, 7(3):348–353, 1974.

Rolf R Mantel. Homothetic preferences and community excess demand functions. *Journal of Economic Theory*, 12(2):197–201, 1976.

Andreu Mas-Colell. The recoverability of consumers' preferences from market demand behavior. *Econometrica*, 45(6):1409–1430, 1977.

Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford Univ. Press, New York, NY, 1995.

Pascal Massart. Some applications of concentration inequalities to statistics. In *Annales de la Faculté des sciences de Toulouse: Mathématiques*, volume 9, pp. 245–303, 2000.

Jiri Matousek (ed.). *Lectures on Discrete Geometry*, volume 212 of *Graduate Texts in Mathematics*. Springer New York, New York, NY, 2002.

T. Matsui. NP-hardness of linear multiplicative programming and related problems. *Journal of Global Optimization*, 9(2):113–119, 1996.

Andreas Maurer. A vector-contraction inequality for rademacher complexities. In *International Conference on Algorithmic Learning Theory (ALT)*. Springer, 2016.

Daniel McFadden, Andreu Mas-Colell, Rolf Mantel, and Marcel K Richter. A characterization of community excess demand functions. *Journal of Economic Theory*, 9(4):361–374, 1974.

Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of Machine Learning*. MIT Press, 2nd edition, 2018.

Yurii Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103 (1):127–152, 2005. doi: 10.1007/s10107-004-0552-5.

Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani (eds.). *Algorithmic Game Theory*. Cambridge University Press, Cambridge, 2007.

Fabrice Talla Nobibon, Laurens Cherchye, Yves Crama, Thomas Demuynck, Bram De Rock, and Frits C. R. Spieksma. Revealed preference tests of collectively rational consumption behavior: Formulations and algorithms. *Operations Research*, 64(6):1197–1216, 2016. URL https://pubsonline.informs.org/doi/abs/10.1287/opre.2016.1527.

Roy Radner. Existence of equilibrium of plans, prices, and price expectations in a sequence of markets. *Econometrica*, 40(2):289–303, 1972.

Paul A. Samuelson. Consumption theory in terms of revealed preference. *Economica*, 15(60): 243–253, 1948.

Wayne Shafer and Hugo Sonnenschein. Chapter 14 Market demand and excess demand functions. In *Handbook of Mathematical Economics*, volume 2, pp. 671–693. Elsevier, 1982.

Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge university press, New York, 2014.

Maurice Sion. On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176, 1958.

Steve Smale. A convergent process of price adjustment and global Newton methods. *Journal of Mathematical Economics*, 3(2):107–120, 1976.

Bart Smeulders, Frits C. R. Spieksma, Laurens Cherchye, and Bram De Rock. Goodness-of-fit measures for revealed preference tests: complexity results and algorithms. *ACM Transactions on Economics and Computation (TEAC)*, 2(1):3:1–3:16, 2014.

Hugo Sonnenschein. The utility hypothesis and market demand theory. *Economic Inquiry*, 11(4): 404–410, 1973a.

Hugo Sonnenschein. Do Walras' identity and continuity characterize the class of community excess demand functions? *Journal of Economic Theory*, 6(4):345–354, 1973b.

R. Srikant. *The Mathematics of Internet Congestion Control*. Systems & Control: Foundations & Applications. Birkhäuser Boston, Boston, MA, 2004.

Kalyan Talluri and Garrett van Ryzin. Revenue management under a general discrete choice model of consumer behavior. *Management Science*, 50(1):15–33, 2004.

Pravin M. Vaidya. A new algorithm for minimizing convex functions over convex sets. *Mathematical Programming*, 73(3):291–341, 1996. ISSN 1436-4646.

Hal Varian. The nonparametric approach to demand analysis. *Econometrica*, 50(4):945–973, 1982.

Hal R. Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1):63–91, 1974.

Hal R. Varian. Goodness-of-fit in optimizing models. *Journal of Econometrics*, 46(1):125–140, 1990.

Roman Vershynin. *High-dimensional Probability: An Introduction with Applications in Data Science*. Number 47 in Cambridge series in statistical and probabilistic mathematics. Cambridge University Press, Cambridge ; New York, NY, 2018.

Léon Walras. *Éléments d'Économie Politique Pure ou Théorie de la Richesse Sociale*. Corbaz/Cie, Lausanne/Paris, 1874.

Yinyu Ye. A path to the Arrow–Debreu competitive market equilibrium. *Mathematical Programming*, 111(1):315–348, 2008.

Morteza Zadimoghaddam and Aaron Roth. Efficiently learning from revealed preference, 2012. URL http://arxiv.org/abs/1211.4150.

Chuwen Zhang, Chang He, Bo Jiang, and Yinyu Ye. The second-order tâtonnement: decentralized interior-point methods for market equilibrium, 2025. URL http://arxiv.org/abs/2508.04822.

ROADMAP OF THE APPENDIX

The complete notations used in this paper are presented below.

Notations. We use the following notations. The notation $\mathbb{R}, \mathbb{Z}, \mathbb{N}$ denotes the set of real, integer, and natural numbers; $\mathbb{R}^n_+$ and $\mathbb{R}^n_{>0}$ denote the sets of n-dimensional vectors with all nonnegative elements, and positive elements, respectively. For a subset $\mathcal{S} \in \mathbb{R}^n$, we use $\mathrm{int}(\mathcal{S})$ to denote the (relative) interior of $\mathcal{S}$. For a natural number $n \in \mathbb{N}$, we let $[n] = \{1, ..., n\}$. We use $[a; b]$ (resp., $[a, b]$) to denote vertical (resp., horizontal) concatenation of arrays or numbers. For a subset $B \subseteq [n]$, we use $\mathbf{x}_B$ (resp., $\mathbf{A}_{BB}$) to denote the subvector (resp., submatrix) of $\mathbf{x}$ (resp., $\mathbf{A}$) obtained by keeping only the elements (resp., rows and columns) indexed by B. We sometimes omit the brackets when there is no confusion. We use $\|\cdot\|$ to denote a proper semi-norm, and $\|\cdot\|_*$ to denote the dual norm; by default, they indicate the induced ℓ_2 norm.

The appendix is organized as follows. Section A reviews the related literature. Sections B and C derive the demand functions, expenditure shares, and primal-dual formulation. Section D proves the generalization guarantees. Section E establishes the convergence of the Frank-Wolfe implementation of CRM with inexact separation oracles and quantifies the effect of additive SEP errors. Sections F and G prove the representation and misspecification results, including the lower-bound construction. Section H analyzes the complexity of the separation problem and develops global and local solution methods. Finally, Section I reports the experimental setup and full numerical results.

CONTENTS OF THE APPENDIX

A Related literature **17**

A.1 Rationalizability and learnability of individual preferences 17

A.2 Rationalization and learning of collective consumption 18

A.3 Fair resource allocation and market equilibrium 18

B Derivation of demand functions and the expenditure shares **19**

B.1 CES consumer 19

B.2 Leontief consumer 20

B.3 Cobb-Douglas consumer 21

B.4 Linear consumer 21

B.5 Homotheticity and the expenditure share 22

C Derivation of the primal and dual problems **22**

D Proof of the performance guarantees for generalization risk **23**

D.1 Preliminaries 23

D.2 Rademacher bounds 25

D.3 Generalization risk for constant wealth (WP_T) 29

D.4 Generalization risk for the extension ($\mathrm{WP}_{+,T}$) 30

E Convergence analysis of the Frank-Wolfe implementation of CRM **32**

E.1 Feasible predictors 32

E.2 Frank-Wolfe with an inexact linear minimization oracle 33

E.3 Application to CRM 34

E.4 Effect of additive SEP errors 36

F Proof of results on misspecification error I **36**

F.1 Transforming real market to lifted market 37

F.2 Verification of Theorem 4.3 38

G Proof of results on misspecification error II **40**

G.1 Two-good economy and total variation 41

G.2 The role of wealth functions 41

G.3 A homothetic agent with unbounded variation 43

G.4 A sharp lower bound 44

H Methods of solving the separation oracle **46**

H.1 Linear case 46

H.2 CES, Cobb-Douglas, and Leontief cases: I, Hardness results 48

H.3 CES, Cobb-Douglas, and Leontief cases: II, An FPTAS for global maximum . . . 51

H.4 CES, Cobb-Douglas, and Leontief cases: III, Spatial branch-and-bound and KKT points 54

I Numerical experiments: setup and full results **55**

I.1 Setup 55

I.2 Ablation study 56

I.3 Prediction under different wealth functions 60

I.4 Computing fair allocations 60

I.5 Performance on approximate PFA 61

A Related literature

We review the most relevant research related to the current study.

A.1 Rationalizability and learnability of individual preferences

Afriat-Varian rationalization has been extensively studied in economics; see, e.g., Chiappori & Rochet (1987); Blundell et al. (2003; 2008; 2015) and the monograph Chambers & Echenique (2016, Chapter 3-5). It is also possible to quantify the extent of "irrationality" of the preferences by the inconsistency index (Varian, 1990; Echenique et al., 2011), but most of these indices cannot be computed very efficiently (Smeulders et al., 2014; Cherchye et al., 2025). To remedy nonsmoothness in Afriat-Varian rationalization, Chiappori & Rochet (1987) show that if the same bundle cannot be induced by different prices, then the data can be rationalized by a smooth, infinitely differentiable, strictly concave utility function. Blundell et al. (2003; 2008; 2015) extends the classical revealed preference approach when the budget is also chosen by the agent, accommodating the effect of wealth changes.

The main drawback of the above methods is that they cannot provide guarantees for out-of-sample prediction. It is also unlikely to be possible universally, since Beigman & Vohra (2006) showed that general demand is not always learnable under the Probably Approximately Correct (PAC) frame-

work[3]. Learnability is generally a feature that only holds for "easy" utilities. For example, Zadimoghaddam & Roth (2012) established polynomial bounds on the sample complexity if the utilities are restricted to be linear and separable concave with Lipschitz gradients. Balcan et al. (2014) provided sample complexity guarantees for a number of important classes, including linear, Leontief, and CES utility functions. A recent line of economic literature has focused on the statistical learnability of general preferences and choices under uncertainty. For example, Basu & Echenique (2020); Chambers et al. (2021) discussed learnable models with finite Vapnik-Chervonenkis (VC) dimensions, in which the choice in the Walrasian setting is one specific application.

A.2 Rationalization and learning of collective consumption

The problem studied in this paper is related to a question asked by Sonnenschein (1973a;b): What restrictions are necessary for a function to be the market excess demand function? This question was answered by Sonnenschein (1973b); Mantel (1974); Debreu (1974); also see Geanakoplos & Polemarchakis (1980); Shafer & Sonnenschein (1982). The Sonnenschein-Mantel-Debreu (SMD) theorem shows that market excess demand can be very general: essentially, any continuous function satisfying homogeneity of degree zero and Walras' law can be the excess demand function of an Arrow-Debreu economy. This generality has two closely related propositions. First, there is little structure in the equilibrium set as well. Indeed, any compact subset of the interior of the unit sphere can arise as the equilibrium set of a well-behaved exchange economy (Mas-Colell, 1977). Second, SMD-type results imply that the convergence of price-adjustment processes (the formal definition is given at (20) shortly below) cannot be guaranteed universally because the excess demand function can almost be arbitrary, while its convergence requires certain regularity conditions (Arrow & Hurwicz, 1958).

Characterization of market demand becomes more difficult when wealth functions are fixed (Shafer & Sonnenschein, 1982). In this setting, the decomposition of aggregate excess demand into a set of individual demand functions is only possible locally (Chiappori & Ekeland, 1999a) around a given price, but not globally (McFadden et al., 1974; Shafer & Sonnenschein, 1982). At the same time, because fixed wealth functions limit the arbitrariness allowed by SMD, the literature has obtained more positive results concerning the stability of the price-adjustment processes.

Several related decomposition problems have also been studied. One line of work generalizes the SMD-type results to the incomplete markets, which is hard due to the Grassmannian from limited market span (Radner, 1972; Brown et al., 1996; Chiappori & Ekeland, 1999b). Another line of literature studies collective consumption lying between purely individual and many-individual market behavior, such as households and firms, with observations of private and public consumption and Lindahl prices (Browning & Chiappori, 1998; Cherchye et al., 2008; 2011). For these models, mixed-integer optimization approaches have been developed in a spirit analogous to Afriat-Varian rationalization.

Finally, our paper is also related to the literature on learning market primitives. Since learning even a single agent's preferences is already difficult, the literature on learning collective consumption is quite limited. To our knowledge, the closest work is Bei et al. (2016), which studies an active-learning setting in which the learner can freely set prices and is therefore not a statistical learning problem.

A.3 Fair resource allocation and market equilibrium

Proportionally fair resource allocation (PFA) and Walrasian market equilibrium (WE) are closely related concepts in economics and computer science (Mas-Colell et al., 1995; Nisan et al., 2007). While finding a fair allocation is generally solvable by convex optimization methods, finding a WE is a fixed-point problem, for which not many cases can be computed efficiently. The two concepts coincide when, for example, the wealth is *fixed* and each agent's utility is homogeneous[4]. In this case, a PFA can be found in polynomial time by maximizing the NSW (Eisenberg & Gale, 1959),

[3] In the same paper, learnable cases could be derived; unfortunately, the PLC utilities cannot be covered. PLC utilities can admit set-valued demand, clearly violating the condition in Beigman & Vohra (2006).

[4] Note that, if the preferences are homothetic, there is no loss of generality in assuming that the utility is homogeneous of degree one.

and the dual variable to the resource constraint in $\mathcal{X}$ is an equilibrium price (Devanur et al., 2002; Ye, 2008). Again, these methods require complete knowledge on the utilities and wealth, which are often difficult to obtain. For the same reason, they do not scale well with the number of agents.

A compelling rationale for pursuing WE is that it allows us to harness the "invisible hand" without requiring complete information about individual agents. One may employ a mechanism that adjusts prices so that the demand matches the supply. The individual demand realized at the resulting equilibrium prices can then serve as a proxy for PFA. Remarkably, even when utilities are not homothetic, every allocation induced by WE achieves at least a 0.69 fraction of the maximum Nash welfare (Garg et al., 2026). Thus, this approach is not merely a theoretical aspiration but can be justified in a broad range of practical settings. As early as the nineteenth century, Walras (1874) introduced the first such mechanism, the tâtonnement process, which works as follows:

$$\mathbf{p} \leftarrow \mathbf{p} + \mathcal{G} \circ \left(\sum_{i \in [m]} \mathbf{x}_i(\mathbf{p}, w_i(\mathbf{p})) - \mathbf{1} \right), \tag{20}$$

where $\mathcal{G} : \mathbb{R}^n \mapsto \mathbb{R}^n$ is some sign-preserving mapping. The moral of this mechanism is very simple: it adjusts the price of good j upward when demand exceeds supply and downward otherwise. The convergence of the tâtonnement and its variants to WE was shown by Arrow & Hurwicz (1958); Smale (1976). In fact, the usage of the tâtonnement is not restricted to homothetic agents and fixed budgets. Worst-case complexity results have been established for linear wealth (Jain, 2007; Ye, 2008; Garg & Végh, 2023), and for utilities satisfying Weak Gross Substitutability (WGS) (Codenotti & Varadarajan, 2007) and bounded elasticity (Goktas et al., 2023), and for markets admitting a convex potential function (Jain & Vazirani, 2010; Cheung et al., 2020), and so on.

While price-adjustment processes do not require full preference revelation, they typically rely on prior structural assumptions on $(\mathscr{X}, \mathscr{W}, \mathscr{U})$ in order to guarantee convergence. For example, a mechanism may assume that agents' utilities satisfy WGS while remaining agnostic about specific parameterizations. In light of SMD, however, such assumptions substantially restrict the class of aggregate demand functions that can arise in a market. Consequently, convergence guarantees derived under these assumptions may not extend to markets encountered in practice.

B Derivation of demand functions and the expenditure shares

We consider the following class of utility functions that are homothetic. In this section, we derive the expenditure shares for these classes. These results will be repetitively used in the our development.

Example 1 (Additive homothetic classes)**.** *Consider the utility function parameterized by* $\mathbf{c} \in \mathbb{R}^n_+$ *and* $r \in [-\infty, 1]$*: $u(\mathbf{x}) = \langle \mathbf{c}, \mathbf{x}^r \rangle^{1/r}$. Depending on r, we have the following component classes: the linear class $\mathscr{H}^{\mathrm{Lin}}$ for $r = 1$, the Leontief case $\mathscr{H}^{\mathrm{Leon}}$ for $r = -\infty$, and the CES class $\mathscr{H}^{\mathrm{CES}}$ for $r \in (-\infty, 1)$, including the Cobb-Douglas class $\mathscr{H}^{\mathrm{Cobbs}}$ if $r = 0$. By letting $\sigma = \frac{r}{1-r}$, these cases can be compactly represented as:*

$$\begin{aligned}
\mathscr{H}^{\mathrm{CES}} &:= \{\boldsymbol{\gamma} : \Delta_n \to \Delta_n \mid (\mathbf{y}, \sigma) \in \mathbb{R}^n \times (-1, \infty),\ \boldsymbol{\gamma}(\mathbf{p}) = \mathrm{softmax}\left(\mathbf{y} - \sigma \log(\mathbf{p})\right)\}, \\
\mathscr{H}^{\mathrm{Cobbs}} &:= \{\boldsymbol{\gamma} : \Delta_n \to \Delta_n \mid \mathbf{y} \in \mathbb{R}^n,\ \boldsymbol{\gamma}(\mathbf{p}) \equiv \boldsymbol{\gamma} = \mathrm{softmax}\left(\mathbf{y}\right)\}, \\
\mathscr{H}^{\mathrm{Leon}} &:= \left\{\boldsymbol{\gamma} : \Delta_n \to \Delta_n \,\middle|\, \mathbf{y} \in \mathbb{R}^n,\ \boldsymbol{\gamma}(\mathbf{p}) = \mathrm{softmax}\left(\mathbf{y} + \log(\mathbf{p})\right)\right\}, \\
\mathscr{H}^{\mathrm{Lin}} &:= \left\{\boldsymbol{\gamma} : \Delta_n \to \Delta_n \,\middle|\, \mathbf{y} \in \mathbb{R}^n,\ \boldsymbol{\gamma}(\mathbf{p}) = \mathbf{e}_{j^*(\mathbf{y},\mathbf{p})},\ j^*(\mathbf{y}, \mathbf{p}) \in \arg\max_{j \in [n]} \{y_j - \log p_j\}\right\}.
\end{aligned}$$

Throughout, the linear class uses the fixed smallest-index selector in (24)*.*

B.1 CES consumer

An agent or android is a said to be a CES consumer if in its utility function, $r \in (-\infty, 1)$. The corresponding demand function is a single-valued mapping with the following analytical form

$$\mathbf{x}(\mathbf{p}) = w \frac{\mathrm{diag}(\mathbf{p})^{-(1+\sigma)} \mathbf{c}^{1+\sigma}}{\langle \mathbf{c}^{1+\sigma}, \mathbf{p}^{-\sigma} \rangle}, \quad \sigma = \frac{r}{1-r} \geq -1.$$

Here the quantity $\sigma + 1 = \frac{1}{1-r}$ is referred to as the elasticity of substitution (Mas-Colell et al., 1995), but a slight abuse of notation, we refer σ as the elasticity for short if no confusion arises. The expenditure share $\boldsymbol{\gamma}(\mathbf{p})$ is summarized as follows.

Fact 1. *For any* $\mathbf{c} \in \mathbb{R}^n_+$, $r \in (-\infty, 1)$, *define the following reparameterization:*

$$\mathbf{y} = (1+\sigma)\log(\mathbf{c}) \in \mathbb{R}^n, \quad \sigma = \frac{r}{1-r} \geq -1.$$

Also define the following mapping $\boldsymbol{\gamma} : \mathcal{P} \mapsto \Delta_n$*:*

$$\boldsymbol{\gamma}(\mathbf{p}) = \frac{\mathbf{c}^{1+\sigma} \circ \mathbf{p}^{-\sigma}}{\langle \mathbf{c}^{1+\sigma}, \mathbf{p}^{-\sigma} \rangle} = \operatorname{softmax}\left(\mathbf{y} - \sigma \log(\mathbf{p})\right), \tag{21}$$

where $\operatorname{softmax} : \mathbb{R}^n \mapsto \Delta_n$ *is the softmax transformation:* $\operatorname{softmax}(\cdot) = \frac{\exp(\cdot)}{\langle \exp(\cdot), \mathbf{1} \rangle}$.

The verification of this fact is straightforward so we leave it to the reader. Motivated from Fact 1, we will use the parameterization $(\mathbf{y}, \sigma)$ in most cases. Scaling the parameters $\mathbf{c}$ (or equivalently $\mathbf{y}$) produces the same expenditure share $\boldsymbol{\gamma}(\mathbf{p})$, viz.

$$\operatorname{softmax}\left(\mathbf{y} - \sigma \log(\mathbf{p})\right) = \operatorname{softmax}\left(\mathbf{y} + a\mathbf{1} - \sigma \log(\mathbf{p})\right), \quad \forall a \in \mathbb{R}. \tag{22}$$

The above shift invariance of the softmax operator is well-known. Now, we can define $\mathscr{H}^{\mathrm{CES}}$ formally.

$$\mathscr{H}^{\mathrm{CES}} := \left\{ \boldsymbol{\gamma} : \mathcal{P} \to \Delta_n \mid (\mathbf{y}, \sigma) \in \mathbb{R}^n \times (-1, \infty),\ \boldsymbol{\gamma}(\mathbf{p}) = \operatorname{softmax}\left(\mathbf{y} - \sigma \log(\mathbf{p})\right) \right\}.$$

The symbol $\mathscr{H}^{\mathrm{CES}}$ recognizes the fact that r only takes values in $(-\infty, 1)$. For the CES case, the output sequence $\{\boldsymbol{\gamma}(\mathbf{p}_k)\}_{k\in[K]}$ is unique given a price sequence $\{\mathbf{p}_k\}_{k\in[K]}$, since any $\boldsymbol{\gamma}$ is a single-valued mapping from $\mathcal{P}$ to Δ_n. Reversely, if given a set of vectors $\{\boldsymbol{\gamma}_1, \ldots, \boldsymbol{\gamma}_K\}$ and each $\boldsymbol{\gamma}_k \in \operatorname{int}(\Delta_n)$, does there exist such a $\boldsymbol{\gamma} \in \mathscr{H}^{\mathrm{CES}}$ that $\boldsymbol{\gamma}_k = \boldsymbol{\gamma}(\mathbf{p}_k)$ holds uniformly? The following fact says the answer is yes if $K = 1$.

Fact 2. *For any* single *pair of price and expenditure share* $(\mathbf{p}, \boldsymbol{\gamma})$, *where* $\mathbf{p} \in \mathcal{P}, \boldsymbol{\gamma} \in \operatorname{int}(\Delta_n)$, *there exists* $(\mathbf{y}, \sigma) \in \mathbb{R}^n \times (-1, \infty)$ *such that* $\boldsymbol{\gamma} = \operatorname{softmax}\left(\mathbf{y} - \sigma \log(\mathbf{p})\right)$.

Proof. Taking log-ratios against $j = 1$:

$$\log(\tfrac{\gamma_j}{\gamma_1}) = (y_j - y_1) - \sigma(\log \tfrac{p_j}{p_1}), \qquad j = 2, \ldots, n. \tag{23}$$

Because of shift invariance (22), we can set $y_1 = 0$ without loss of generality. This gives $n-1$ equations in n unknowns $(y_2, \ldots, y_n, \sigma)$, leaving σ as a free parameter. For any fixed $\sigma > -1$, we recover

$$y_j = \log(\tfrac{\gamma_j}{\gamma_1}) + \sigma(\log \tfrac{p_j}{p_1}), \qquad j = 2, \ldots, n.$$

The proof is complete. ■

In the meantime, with $K > 1$ data points, a linear system analogous to (23) is overdetermined. It is not possible to recover the parameters reliably whenever there are more than two different samples.

B.2 Leontief consumer

In the Leontief case, the utility function is $u(\mathbf{x}) = \min_{j\in[n]} \frac{x_j}{c_j}, \mathbf{c} \in \mathbb{R}^n_{>0}$, and the corresponding UMP admits the perfect-complements demand $\mathbf{x}(\mathbf{p}) = \frac{w\mathbf{c}}{\langle \mathbf{p}, \mathbf{c} \rangle}$. We present the following fact for $\boldsymbol{\gamma}$ in parallel to Fact 1.

Fact 3. *For any* $\mathbf{c} \in \mathbb{R}^n_{>0}$, *set the reparameterization* $\mathbf{y} = \log(\mathbf{c}) \in \mathbb{R}^n$. *Then the expenditure share of a Leontief consumer at price* $\mathbf{p}$ *is* $\boldsymbol{\gamma}(\mathbf{p}) = \frac{\operatorname{diag}(\mathbf{p})\mathbf{c}}{\langle \mathbf{p}, \mathbf{c} \rangle} = \operatorname{softmax}\left(\mathbf{y} + \log(\mathbf{p})\right)$.

The right-hand side recovers the CES expression (21) at the boundary $\sigma = -1$, where the substitution effect vanishes. Similar to the CES case, $\mathscr{H}^{\mathrm{Leon}}$ can be written as:

$$\mathscr{H}^{\mathrm{Leon}} := \left\{ \boldsymbol{\gamma} : \mathcal{P} \to \Delta_n \;\middle|\; \mathbf{y} \in \mathbb{R}^n,\ \boldsymbol{\gamma}(\mathbf{p}) = \operatorname{softmax}\left(\mathbf{y} + \log(\mathbf{p})\right) \right\}.$$

Analogously, a single CES observation is always representable.

Fact 4. *For any single pair* $(\mathbf{p}, \boldsymbol{\gamma})$ *with* $\mathbf{p} \in \mathbb{R}^n_{>0}$ *and* $\boldsymbol{\gamma} \in \operatorname{int}(\Delta_n)$, *there exists* $\mathbf{y} \in \mathbb{R}^n$ *such that* $\boldsymbol{\gamma} = \operatorname{softmax}\left(\mathbf{y} + \log(\mathbf{p})\right)$.

The proof is similar to that for Fact 2.

B.3 COBB-DOUGLAS CONSUMER

The Cobb-Douglas consumer is the homothetic case $r \to 0$, equivalently the unit-elasticity case $\sigma = 0$ of Fact 1. Its utility is $u(\mathbf{x}) = \prod_{j=1}^{n} x_j^{c_j}, \mathbf{c} \in \mathbb{R}^n_{>0}$, and the corresponding UMP admits the demand $\mathbf{x}(\mathbf{p}) = \frac{w}{\langle \mathbf{c}, \mathbf{1} \rangle} \mathrm{diag}(\mathbf{p})^{-1}\mathbf{c}$. We have the following immediately,

$$\mathscr{H}^{\mathrm{Cobbs}} := \{\boldsymbol{\gamma} : \mathcal{P} \to \Delta_n \mid \mathbf{y} \in \mathbb{R}^n,\ \boldsymbol{\gamma}(\mathbf{p}) \equiv \boldsymbol{\gamma} = \mathrm{softmax}\,(\mathbf{y})\}\,.$$

B.4 LINEAR CONSUMER

For the linear case, its utility is given by $u(\mathbf{x}) = \langle \mathbf{c}, \mathbf{x} \rangle$, with $\mathbf{c} \in \mathbb{R}^n_+$. The demand concentrates expenditure on the smallest-index bang-per-buck maximizer:

$$j^* := \min \arg\max_{j \in [n]} \frac{c_j}{p_j}, \qquad \mathbf{x}(\mathbf{p}) = \frac{w}{p_{j^*}} \mathbf{e}_{j^*}. \tag{24}$$

Fact 5. *For any $\mathbf{c} \in \mathbb{R}^n_{>0}$, set $\mathbf{y} = \log(\mathbf{c}) \in \mathbb{R}^n$. The expenditure share of a linear consumer at price $\mathbf{p}$ is*

$$\boldsymbol{\gamma}(\mathbf{p}) = \mathbf{e}_{j^*(\mathbf{y},\mathbf{p})}, \qquad j^*(\mathbf{y}, \mathbf{p}) := \min \arg\max_{j \in [n]} \{y_j - \log p_j\}.$$

The same priority order applies at every price, so identical prices have identical expenditure shares. The linear class is therefore

$$\mathscr{H}^{\mathrm{Lin}} := \left\{ \boldsymbol{\gamma} : \mathcal{P} \to \Delta_n \;\middle|\; \mathbf{y} \in \mathbb{R}^n,\ \boldsymbol{\gamma}(\mathbf{p}) = \mathbf{e}_{j^*(\mathbf{y},\mathbf{p})},\ j^*(\mathbf{y}, \mathbf{p}) = \min \arg\max_{j \in [n]} \{y_j - \log p_j\} \right\}.$$

Fact 6. *For any single pair $(\mathbf{p}, \mathbf{e}_{j^*})$ with $\mathbf{p} \in \mathbb{R}^n_{>0}$ and $j^* \in [n]$, there exists $\boldsymbol{\gamma} \in \mathscr{H}^{\mathrm{Lin}}$ such that $\boldsymbol{\gamma}(\mathbf{p}) = \mathbf{e}_{j^*}$.*

Proof. Choose

$$y_{j^*} = 0, \qquad y_j = \log\left(\frac{p_j}{p_{j^*}}\right) - 1, \quad j \neq j^*.$$

Then $y_{j^*} - \log p_{j^*} = y_j - \log p_j + 1$ for every $j \neq j^*$, so j^* is the unique maximizer. ■

For K samples with observed shares $\boldsymbol{\gamma}_k = \mathbf{e}_{j_k}$, compatibility with this fixed selector requires

$$\begin{aligned} y_{j_k} - y_\ell &> \log\left(\frac{p_{k,j_k}}{p_{k,\ell}}\right), \quad \ell < j_k,\ k \in [K], \\ y_{j_k} - y_\ell &\geq \log\left(\frac{p_{k,j_k}}{p_{k,\ell}}\right), \quad \ell > j_k,\ k \in [K], \\ y_n &= 0. \end{aligned} \tag{25}$$

The strict comparisons exclude ties with higher-priority goods. For a choice pattern $\mathbf{j} = (j_1, \ldots, j_K)$, define its maximum common margin by the linear program

$$\begin{aligned} \eta^*(\mathbf{j}) = \max_{\mathbf{y}, \eta}\ & \eta \\ \text{s.t.}\ \ & y_{j_k} - y_\ell \geq \log\left(\frac{p_{k,j_k}}{p_{k,\ell}}\right) + \eta, \quad \ell < j_k,\ k \in [K], \\ & y_{j_k} - y_\ell \geq \log\left(\frac{p_{k,j_k}}{p_{k,\ell}}\right), \qquad \ell > j_k,\ k \in [K], \\ & y_n = 0, \qquad 0 \leq \eta \leq 1. \end{aligned} \tag{26}$$

Set $\eta^*(\mathbf{j}) = -\infty$ when this program is infeasible. The pattern is realizable if and only if $\eta^*(\mathbf{j}) > 0$: a positive margin implies (25), and any solution of that system has a positive minimum over its finitely many strict gaps. When there are no strict comparisons, weak feasibility implies $\eta^*(\mathbf{j}) = 1$. This test optimizes the margin rather than prescribing a fixed positive tolerance. It applies to the unrestricted parameterization of $\mathscr{H}^{\mathrm{Lin}}$ above and is used in the separation oracle in Section H.1.

B.5 Homotheticity and the expenditure share

Lemma B.1. *The expenditure share $\boldsymbol{\gamma}(\mathbf{p}, w)$ is independent of w iff the preferences are homothetic.*

Proof. Suppose u is homothetic, then there exists a strictly increasing function ϕ and a homogeneous of degree one function $\tilde{u}$ such that $u = \phi \circ \tilde{u}$. Substituting $\mathbf{y} = \frac{1}{\lambda}\mathbf{x}$ in the UMP at $(\mathbf{p}, \lambda w)$ and using the homogeneity of $\tilde{u}$,

$$\mathbf{x}(\mathbf{p}, \lambda w) = \arg \max_{\langle \mathbf{p}, \mathbf{x} \rangle \leq \lambda w} u(\mathbf{x}) = \lambda \arg \max_{\langle \mathbf{p}, \mathbf{y} \rangle \leq w} \phi(\lambda \tilde{u}(\mathbf{y})).$$

Both $\phi(\lambda \tilde{u}(\cdot))$ and $u(\cdot)$ are strictly increasing transforms of $\tilde{u}(\cdot)$, so the inner argmax is $\mathbf{x}(\mathbf{p}, w)$. Setting $\lambda = \frac{1}{w}$ gives $\mathbf{x}(\mathbf{p}, w) = w\mathbf{x}(\mathbf{p}, 1)$, and

$$\boldsymbol{\gamma}(\mathbf{p}, w) = \operatorname{diag}(\mathbf{p})\mathbf{x}(\mathbf{p}, 1),$$

independent of w. Conversely, suppose $\boldsymbol{\gamma}(\mathbf{p}, w) \equiv \boldsymbol{\gamma}(\mathbf{p})$, equivalently the Walrasian demand is homogeneous of degree one in wealth,

$$\mathbf{x}(\mathbf{p}, w) = w\, \mathbf{x}(\mathbf{p}, 1), \qquad \forall \mathbf{p} \in \mathcal{P},\ w > 0,$$

so every wealth-expansion path $\{\mathbf{x}(\mathbf{p}, w) : w > 0\}$ is a ray through the origin. For continuous, locally non-satiated, and convex preferences (so that $\mathbf{x}(\mathbf{p}, w)$ is well defined), this is precisely the standard characterization of homotheticity: the Walrasian demand is homogeneous of degree one in wealth at every price if and only if $\succeq$ is homothetic (Mas-Colell et al., 1995, Chapter 3). Hence the preferences are homothetic. ∎

C Derivation of the primal and dual problems

Rewrite the semi-infinite primal (WP):

$$\begin{aligned} L^* = \inf_{\omega \in \mathfrak{M}(\mathscr{H}), \mathbf{s}_1, \ldots, \mathbf{s}_K} \ & \tfrac{1}{K} \sum_{k=1}^{K} \|\mathbf{s}_k\|, \\ \text{s.t.} \ \ & \mathbf{s}_k + \int_{\mathscr{H}} \boldsymbol{\gamma}(\mathbf{p}_k)\, \mathrm{d}\omega(\boldsymbol{\gamma}) = \mathbf{g}_k, \forall k \\ & \mathbf{s}_k \in \mathbb{R}^n, \forall k. \end{aligned}$$

Write $\mathbf{u}_k = \mathbf{U}\mathbf{e}_k \in \mathbb{R}^n$ for each sample k and $\mu \in \mathbb{R}$ for the simplex constraint $\int_{\mathscr{H}} \mathrm{d}\omega = 1$. The Lagrangian is

$$\begin{aligned} \mathcal{L}(\omega, \mathbf{s}; \mathbf{u}, \mu) = \tfrac{1}{K} \Bigg[& \sum_k \big(\|\mathbf{s}_k\| - \langle \mathbf{u}_k, \mathbf{s}_k \rangle \big) \\ & + \int_{\mathscr{H}} \Big(\mu - \sum_k \langle \mathbf{u}_k, \boldsymbol{\gamma}(\mathbf{p}_k) \rangle \Big) \mathrm{d}\omega(\boldsymbol{\gamma}) \\ & + \sum_k \langle \mathbf{u}_k, \mathbf{g}_k \rangle - \mu \Bigg]. \end{aligned} \tag{27}$$

After dualizing the unit-mass constraint, the Lagrangian is minimized over finite nonnegative measures ω. For a semi-norm, the notation $\|\mathbf{u}\|_* \leq 1$ denotes membership in its compact convex polar ball:

$$\mathcal{B}_* := \{\mathbf{u} \in \mathbb{R}^n : \langle \mathbf{u}, \mathbf{v} \rangle \leq \|\mathbf{v}\|, \quad \forall \mathbf{v} \in \mathbb{R}^n\}, \qquad \|\mathbf{v}\| = \max_{\mathbf{u} \in \mathcal{B}_*} \langle \mathbf{u}, \mathbf{v} \rangle.$$

Note $\inf_{\mathbf{s}_k} \big[\|\mathbf{s}_k\| - \langle \mathbf{u}_k, \mathbf{s}_k \rangle \big] > -\infty$ iff $\|\mathbf{u}_k\|_* \leq 1$, so we can write the dual problem

$$\begin{aligned} L^* \geq \max_{\mathbf{u}, \mu} \ & \tfrac{1}{K} \Big(\sum_k \langle \mathbf{u}_k, \mathbf{g}_k \rangle - \mu \Big) \\ \text{s.t.} \ \ & \sum_{k=1}^{K} \langle \mathbf{u}_k, \boldsymbol{\gamma}(\mathbf{p}_k) \rangle \leq \mu, \quad \forall \boldsymbol{\gamma} \in \mathscr{H}, \\ & \|\mathbf{u}_k\|_* \leq 1, \quad \forall k \in [K], \end{aligned}$$

which is exactly (WD).

Lemma C.1 (Strong duality)**.** *Suppose $\mathscr{H}$ is a nonempty class of homothetic expenditure-share functions with Borel measurable evaluation maps. The infimum in* (WP) *equals the maximum in* (WD)*, and the dual maximum is attained.*

Proof. Define the sample evaluations and their closed convex hull by

$$\mathbf{A}(\boldsymbol{\gamma}) := [\boldsymbol{\gamma}(\mathbf{p}_1) \quad \cdots \quad \boldsymbol{\gamma}(\mathbf{p}_K)], \qquad \mathcal{S} := \{\mathbf{A}(\boldsymbol{\gamma}) : \boldsymbol{\gamma} \in \mathscr{H}\}, \qquad \mathcal{C} := \overline{\mathrm{conv}}(\mathcal{S}).$$

For a matrix of predicted expenditure shares, with one column per observed price, we can rewrite the empirical loss by

$$F(\mathbf{X}) := \frac{1}{K}\sum_{k=1}^{K} \|\mathbf{g}_k - \mathbf{X}\mathbf{e}_k\|, \qquad \mathbf{X} \in \mathbb{R}^{n\times K}.$$

Since each column of $\mathbf{A}(\boldsymbol{\gamma})$ belongs to Δ_n, the set $\mathcal{C}$ is compact and convex. Every probability measure has its evaluation barycenter in $\mathcal{C}$, while finitely supported measures realize every point of $\mathrm{conv}(\mathcal{S})$. Since F is continuous and $\mathrm{conv}(\mathcal{S})$ is dense in $\mathcal{C}$, it follows that

$$L^* = \inf_{\omega\in\mathfrak{M}(\mathscr{H})} F\left(\int_{\mathscr{H}} \mathbf{A}(\boldsymbol{\gamma})\mathrm{d}\omega\right) = \inf_{\mathbf{X}\in\mathrm{conv}(\mathcal{S})} F(\mathbf{X}) = \min_{\mathbf{X}\in\mathcal{C}} F(\mathbf{X}).$$

Define the compact convex dual domain

$$\mathcal{U} := \{\mathbf{U} \in \mathbb{R}^{n\times K} : \mathbf{U}\mathbf{e}_k \in \mathcal{B}_*, \quad k \in [K]\}.$$

The polar representation of the loss and the minimax theorem yield

$$\begin{aligned} L^* &= \min_{\mathbf{X}\in\mathcal{C}} \max_{\mathbf{U}\in\mathcal{U}} \frac{1}{K}\sum_{k=1}^{K} \langle \mathbf{U}\mathbf{e}_k, \mathbf{g}_k - \mathbf{X}\mathbf{e}_k\rangle \\ &\overset{(a)}{=} \max_{\mathbf{U}\in\mathcal{U}} \min_{\mathbf{X}\in\mathcal{C}} \frac{1}{K}\sum_{k=1}^{K} \langle \mathbf{U}\mathbf{e}_k, \mathbf{g}_k - \mathbf{X}\mathbf{e}_k\rangle \\ &= \max_{\mathbf{U}\in\mathcal{U}} \frac{1}{K}\left(\sum_{k=1}^{K} \langle \mathbf{U}\mathbf{e}_k, \mathbf{g}_k\rangle - \sup_{\boldsymbol{\gamma}\in\mathscr{H}} \sum_{k=1}^{K} \langle \mathbf{U}\mathbf{e}_k, \boldsymbol{\gamma}(\mathbf{p}_k)\rangle\right), \end{aligned}$$

where (a) follows from Sion (1958, Theorem 3.4), since the objective is continuous and affine in each variable on compact convex domains. The support function of $\mathcal{S}$ is finite and continuous because $\mathcal{S}$ is bounded, so the last maximum is attained on $\mathcal{U}$. Taking μ equal to the displayed supremum yields exactly (WD) and an attaining dual pair. ■

If $\mathcal{S}$ is compact, then $\mathcal{C} = \mathrm{conv}(\mathcal{S})$, and Carathéodory's theorem yields a finitely supported primal optimizer. Without this additional condition, the primal infimum need not be attained; the finite master problem (WP_T) still attains its minimum on Δ_T.

D Proof of the performance guarantees for generalization risk

We prove Theorem 4.1 in three steps:

(1) We first establish the Lipschitz-contraction and Rademacher-Gaussian comparison results used in the analysis.
(2) Next, we bound the empirical Rademacher complexity of a finite union of component classes and specialize the bounds to the utility classes considered in this paper.
(3) Finally, we prove the sharper constant-wealth guarantee in Theorem D.1 and then extend the argument to linear wealth and bounded extra supply in Section D.4, yielding Theorem 4.1.

D.1 Preliminaries

We begin with the auxiliary contraction and comparison results.

Lemma D.1 (Lipschitz continuity of the loss function)**.** *For any fixed $\mathbf{t} \in \mathbb{R}^n$, the loss $\ell(\cdot, \mathbf{t})$ is* 1*-Lipschitz in the first argument:*

$$|\ell(\mathbf{x}, \mathbf{t}) - \ell(\mathbf{x}', \mathbf{t})| \le \|\mathbf{x} - \mathbf{x}'\|_\infty \le \|\mathbf{x} - \mathbf{x}'\|_2, \quad \forall \mathbf{x}, \mathbf{x}' \in \mathbb{R}^n. \tag{28}$$

The proof is straightforward. The following lemma (Maurer, 2016) extends standard Lipschitz contraction inequality (Ledoux & Talagrand, 1991).

Lemma D.2 (Vector-contraction inequality)**.** *Let $\mathscr{H}$ be a class of functions from an input space $\mathcal{W}$ to $\mathbb{R}^n$. For any $k = 1, \ldots, K$, let $\Phi_k : \mathbb{R}^n \mapsto \mathbb{R}$ be M_0-Lipschitz continuous function. For any sample $\boldsymbol{\omega}_1, \ldots, \boldsymbol{\omega}_K \in \mathcal{W}$ of size K, the following inequality holds:*

$$\frac{1}{K}\mathbb{E}_{\varepsilon}\left[\sup_{\boldsymbol{\gamma}\in\mathscr{H}}\sum_{k=1}^{K}\varepsilon_k\Phi_k(\boldsymbol{\gamma}(\boldsymbol{\omega}_k))\right] \leq \frac{\sqrt{2}M_0}{K}\mathbb{E}_{\boldsymbol{\xi}}\left[\sup_{\boldsymbol{\gamma}\in\mathscr{H}}\sum_{k=1}^{K}\langle\boldsymbol{\xi}_k, \boldsymbol{\gamma}(\boldsymbol{\omega}_k)\rangle\right].$$

where $\boldsymbol{\varepsilon} = [\varepsilon_k]$ and $\boldsymbol{\xi}_k \in [\xi_{k,n}], \forall k \in [K]$ are independent Rademacher variables over $\{\pm 1\}$.

Specializing $\boldsymbol{\omega}_k = \mathbf{p}_k$, and $\Phi_k(\cdot) = \ell(\cdot, \mathbf{g}_k)$, we can write the above result using short-hand notations (Maurer, 2016) for the empirical Rademacher complexity of $\mathscr{H}$ and of the loss class $\ell \circ \mathscr{H}$ as

$$\hat{\mathcal{R}}_K(\mathscr{H}) := \frac{1}{K}\mathbb{E}_{\boldsymbol{\xi}}\left[\sup_{\boldsymbol{\gamma}\in\mathscr{H}}\sum_{k=1}^{K}\langle\boldsymbol{\xi}_k, \boldsymbol{\gamma}(\mathbf{p}_k)\rangle\right],$$

$$\hat{\mathcal{R}}_K(\ell \circ \mathscr{H}) := \frac{1}{K}\mathbb{E}_{\varepsilon}\left[\sup_{\boldsymbol{\gamma}\in\mathscr{H}}\sum_{k=1}^{K}\varepsilon_k\ell(\boldsymbol{\gamma}(\mathbf{p}_k), \mathbf{g}_k)\right], \tag{29}$$

where $\varepsilon_k \in \{\pm 1\}$ is the scalar Rademacher variable. Lemma D.2 can be restated, with a slight abuse of notations, as follows:

$$\hat{\mathcal{R}}_K(\ell \circ \mathscr{H}) \leq \sqrt{2} \cdot \hat{\mathcal{R}}_K(\mathscr{H}). \tag{30}$$

Since for any k, Φ_k is 1-Lipschitz due to Lemma D.1. We will use this inequality in the sequel. Before that, let us bound the empirical Rademacher complexity of $\mathscr{H}$, which depends on the component classes. We will need the following comparison lemma for Rademacher and Gaussian averages.

Lemma D.3 (Comparison lemma)**.** *Let $\boldsymbol{\xi} \in \{\pm 1\}^n$ be a Rademacher vector with i.i.d. coordinates and $\boldsymbol{\zeta} \sim \mathbf{N}(\mathbf{0}, \mathbf{I}_n)$ a standard Gaussian vector. For any set $\mathcal{V} \subseteq \mathbb{R}^n$,*

$$\mathbb{E}_{\boldsymbol{\xi}}\left[\sup_{\mathbf{v}\in\mathcal{V}}\langle\boldsymbol{\xi}, \mathbf{v}\rangle\right] \leq \sqrt{\tfrac{\pi}{2}} \cdot \mathbb{E}_{\boldsymbol{\zeta}}\left[\sup_{\mathbf{v}\in\mathcal{V}}\langle\boldsymbol{\zeta}, \mathbf{v}\rangle\right]. \tag{31}$$

Proof. We use the fact that by a Rademacher variable ε_j, $\varepsilon_j|\zeta_j| \sim \mathbf{N}(0, 1)$, then,

$$\mathbb{E}|\zeta_j| = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}|\zeta|\exp\{-\zeta^2/2\}d\zeta = \frac{2}{\sqrt{2\pi}}\int_{0}^{\infty}\zeta\exp\{-\zeta^2/2\}d\zeta = \sqrt{\frac{2}{\pi}}. \tag{32}$$

Thus, we have

$$\begin{aligned}
\mathbb{E}_{\boldsymbol{\zeta}}\left[\sup_{\mathbf{v}\in\mathcal{V}}\langle\boldsymbol{\zeta}, \mathbf{v}\rangle\right] &= \mathbb{E}_{\boldsymbol{\zeta}}\left[\sup_{\mathbf{v}\in\mathcal{V}}\sum_{j=1}^{n}\zeta_j v_j\right] \\
&= \mathbb{E}_{\boldsymbol{\varepsilon}}\left[\mathbb{E}_{\boldsymbol{\zeta}}\left[\sup_{\mathbf{v}\in\mathcal{V}}\sum_{j=1}^{n}\varepsilon_j|\zeta_j| \cdot v_j\right] \middle| \boldsymbol{\varepsilon}\right] \\
&\underset{\square_1}{\geq} \mathbb{E}_{\boldsymbol{\varepsilon}}\left[\sup_{\mathbf{v}\in\mathcal{V}}\sum_{j=1}^{n}\mathbb{E}[|\zeta_j|]\varepsilon_j v_j\right] \\
&\underset{(32)}{=} \sqrt{\tfrac{2}{\pi}}\mathbb{E}_{\boldsymbol{\varepsilon}}\sup_{\mathbf{v}\in\mathcal{V}}\sum_{j=1}^{n}\varepsilon_j v_j = \sqrt{\tfrac{2}{\pi}} \cdot \mathbb{E}_{\boldsymbol{\xi}}\sup_{\mathbf{v}\in\mathcal{V}}\langle\boldsymbol{\xi}, \mathbf{v}\rangle,
\end{aligned}$$

where $\square_1$ due to Jensen's inequality applied to the inner expectation and the convex supremum. Rearranging the terms gives (31). This completes the proof. ■

The above lemma is well-known in the literature; see, e.g., Ledoux & Talagrand (1991); Bartlett & Mendelson (2003). Perhaps an explicit specification of the constant $\sqrt{\frac{\pi}{2}}$, which is dimension-free,

is less known in the literature (see a course (Duchi, 2017) for more details), so we include it here for the completeness of presentation.

D.2 Rademacher bounds

As we use a finite family of $\bar{f}$ component classes to construct $\mathscr{H}$, by definition, $\mathscr{H} = \bigcup_{f\in[\bar{f}]} \mathscr{H}^f$, we have the following lemma.

Lemma D.4. *For any finite union $\mathscr{H} = \bigcup_{f\in[\bar{f}]} \mathscr{H}^f$, it holds that*

$$\hat{\mathcal{R}}_K(\mathscr{H}) \le \max_{f\in[\bar{f}]} \hat{\mathcal{R}}_K(\mathscr{H}^f) + \sqrt{\frac{2\log \bar{f}}{K}}. \tag{33}$$

Proof. Define

$$Z_f(\boldsymbol{\xi}) := \sup_{\boldsymbol{\gamma}\in\mathscr{H}^f} \frac{1}{K}\sum_{k=1}^{K} \langle \boldsymbol{\xi}_k, \boldsymbol{\gamma}(\mathbf{p}_k)\rangle, \quad Z(\boldsymbol{\xi}) := \max_{f\in[\bar{f}]} Z_f(\boldsymbol{\xi}),$$

so that $\hat{\mathcal{R}}_K(\mathscr{H}^f) = \mathbb{E}_{\boldsymbol{\xi}}[Z_f]$ and $\hat{\mathcal{R}}_K(\mathscr{H}) = \mathbb{E}_{\boldsymbol{\xi}}[Z]$. Fix any sample index $k \in [K]$, and let $\boldsymbol{\xi}'$ agree with $\boldsymbol{\xi}$ except the k-th block $\boldsymbol{\xi}_k$. Then

$$\begin{aligned} Z_f(\boldsymbol{\xi}') - Z_f(\boldsymbol{\xi}) &= \sup_{\boldsymbol{\gamma}\in\mathscr{H}^f} \frac{1}{K}\sum_{k=1}^{K} \langle \boldsymbol{\xi}'_k, \boldsymbol{\gamma}(\mathbf{p}_k)\rangle - \sup_{\boldsymbol{\gamma}\in\mathscr{H}^f} \frac{1}{K}\sum_{k=1}^{K} \langle \boldsymbol{\xi}_k, \boldsymbol{\gamma}(\mathbf{p}_k)\rangle \\ &\underset{(\square_1)}{\le} \sup_{\boldsymbol{\gamma}\in\mathscr{H}^f} \tfrac{1}{K}\langle \boldsymbol{\xi}'_k - \boldsymbol{\xi}_k, \boldsymbol{\gamma}(\mathbf{p}_k)\rangle \\ &\le \sup_{\boldsymbol{\gamma}\in\mathscr{H}^f} \tfrac{1}{K}\|\boldsymbol{\xi}'_k - \boldsymbol{\xi}_k\|_\infty \cdot \|\boldsymbol{\gamma}(\mathbf{p}_k)\|_1 \underset{(\square_2)}{\le} \tfrac{2}{K}, \end{aligned}$$

where $(\square_1)$ uses $\sup F' - \sup F \le \sup(F' - F)$, and $(\square_2)$ is due to $\|\boldsymbol{\xi}'_k - \boldsymbol{\xi}_k\|_\infty \le 2$ and $\|\boldsymbol{\gamma}(\mathbf{p}_k)\|_1 = 1$. Exchanging $\boldsymbol{\xi}$ and $\boldsymbol{\xi}'$ gives the reverse inequality $Z_f(\boldsymbol{\xi}) - Z_f(\boldsymbol{\xi}') \le \frac{2}{K}$. Thus, Z_f has bounded differences with constants $c_k = \frac{2}{K}$. By McDiarmid's inequality (e.g., Vershynin (2018, Theorem 2.9.1)),

$$\mathbb{P}_{\boldsymbol{\xi}}[Z_f(\boldsymbol{\xi}) - \mathbb{E}_{\boldsymbol{\xi}} Z_f \ge t] \le \exp\left(\frac{-2t^2}{\sum_{k=1}^{K} c_k^2}\right) = \exp\left(\frac{-t^2}{2/K}\right);$$

that is, $Z_f - \mathbb{E}_{\boldsymbol{\xi}} Z_f$ is sub-Gaussian. Hence,

$$\mathbb{E}_{\boldsymbol{\xi}}\left[\exp\left(t(Z_f - \mathbb{E}_{\boldsymbol{\xi}} Z_f)\right)\right] \le \exp\left(\frac{t^2}{2K}\right), \quad \forall t \in \mathbb{R}. \tag{34}$$

Therefore, for any $t > 0$,

$$\begin{aligned} \mathbb{E}_{\boldsymbol{\xi}}\left[\max_{f\in[\bar{f}]}(Z_f - \mathbb{E}_{\boldsymbol{\xi}} Z_f)\right] &\underset{(\square_3)}{\le} \frac{1}{t}\log\left(\mathbb{E}_{\boldsymbol{\xi}}\left[\exp\left(t\max_{f\in[\bar{f}]}(Z_f - \mathbb{E}_{\boldsymbol{\xi}} Z_f)\right)\right]\right) \\ &\underset{(\square_4)}{\le} \frac{1}{t}\log\left(\mathbb{E}_{\boldsymbol{\xi}}\left[\sum_{f\in[\bar{f}]}\exp\left(t(Z_f - \mathbb{E}_{\boldsymbol{\xi}} Z_f)\right)\right]\right) \\ &\underset{(34)}{\le} \frac{1}{t}\log\left(\exp(t^2/2K)\bar{f}\right) \le \frac{\log \bar{f}}{t} + \frac{t}{2K}. \end{aligned}$$

For $(\square_3)$, we note that by taking $W = \max_{f\in[\bar{f}]}(Z_f - \mathbb{E}_{\boldsymbol{\xi}} Z_f)$, $e^{t\mathbb{E}_{\boldsymbol{\xi}}[W]} \le \mathbb{E}_{\boldsymbol{\xi}}[e^{tW}]$ because of Jensen's inequality; taking logarithms and dividing by t gives $\mathbb{E}_{\boldsymbol{\xi}}[W] \le \frac{1}{t}\log \mathbb{E}_{\boldsymbol{\xi}}[e^{tW}]$ as claimed. Equation $(\square_4)$ is because each term is nonnegative. Optimizing the right-hand side over $t > 0$ by taking $t = \sqrt{2K\log \bar{f}}$ yields

$$\mathbb{E}_{\boldsymbol{\xi}}\left[\max_{f\in[\bar{f}]}(Z_f - \mathbb{E}_{\boldsymbol{\xi}} Z_f)\right] \le \sqrt{\frac{2\log \bar{f}}{K}}.$$

Finally,

$$\hat{\mathcal{R}}_K(\mathscr{H}) = \mathbb{E}_{\boldsymbol{\xi}}\left[\max_{f\in[\bar{f}]} Z_f\right] \le \max_{f\in[\bar{f}]} \mathbb{E}_{\boldsymbol{\xi}}\left[Z_f\right] + \mathbb{E}_{\boldsymbol{\xi}}\left[\max_{f\in[\bar{f}]}(Z_f - \mathbb{E}_{\boldsymbol{\xi}} Z_f)\right]$$
$$\le \max_{f\in[\bar{f}]} \hat{\mathcal{R}}_K(\mathscr{H}^f) + \sqrt{\frac{2\log \bar{f}}{K}}.$$

This proves (33). ■

Since $\bar{f}$ is fixed (e.g., $\bar{f} = 3$ for $\mathscr{H} = \mathscr{H}^{\mathrm{CES}} \cup \mathscr{H}^{\mathrm{Leon}} \cup \mathscr{H}^{\mathrm{Lin}}$), the union penalty $\sqrt{2\log \bar{f}/K}$ is $O(K^{-1/2})$ — the same rate as the per-class Rademacher bound established next, and absorbed into its $\mathrm{poly}/\sqrt{K}$ envelope. It therefore suffices to verify, for each component class, that

$$\boxed{\hat{\mathcal{R}}_K(\mathscr{H}^f) \le \tilde{\mathcal{O}}\left(\frac{1}{\sqrt{K}}\right).} \tag{35}$$

D.2.1 CES, Cobb-Douglas, and Leontief classes

We make the following assumptions for these classes.

Assumption D.1 (Boundedness)**.** *There exist finite constants $D_{\mathbf{y}} \ge 0, D_\sigma \ge 1, D_{\mathbf{p}} \ge 0$. For any $k \in [K]$, $\|\log(\mathbf{p}_k)\|_2 \le D_{\mathbf{p}}$. For any $\boldsymbol{\gamma} \in \mathscr{H}^{\mathrm{CES}} \cup \mathscr{H}^{\mathrm{Leon}}$, parameter σ satisfies $\sigma \in [-1, D_\sigma]$, parameter $\mathbf{y}$ is normalized by $\langle \mathbf{1}, \mathbf{y}\rangle = 0$ and satisfies $\|\mathbf{y}\|_2 \le D_{\mathbf{y}}$.*

Note the range for σ implies $-\infty \le r \le \frac{D_\sigma}{1+D_\sigma} < 1$ or $r = 1$. Shift invariance removes only the common-offset degree of freedom and permits the normalization $\langle \mathbf{1}, \mathbf{y}\rangle = 0$ without loss of generality. It does not bound relative coordinates: under Assumption D.1, $|y_i - y_j| \le \sqrt{2} D_{\mathbf{y}}$, so increasingly concentrated expenditure shares can require a larger $D_{\mathbf{y}}$ when prices and σ are fixed. Thus, the finite bound $D_{\mathbf{y}}$ is a modeling restriction, and the result below applies to this bounded parameter class. We need the following lemma.

Lemma D.5 (Lipschitz continuity of the softmax operator)**.** *The softmax function is $\frac{1}{2}$-Lipschitz continuous:*

$$\|\mathrm{softmax}\,(\mathbf{x}) - \mathrm{softmax}\,(\mathbf{x}')\,\|_2 \le \tfrac{1}{2}\|\mathbf{x} - \mathbf{x}'\|_2, \quad \forall \mathbf{x}, \mathbf{x}' \in \mathbb{R}^n. \tag{36}$$

Proof. To show the result above, it suffices to bound the operator norm of the Jacobian mapping. Let $\mathbf{s} = \mathrm{softmax}\,(\mathbf{x})$. Then $\nabla \mathrm{softmax}\,(\mathbf{x}) = \mathrm{diag}(\mathbf{s}) - \mathbf{s}\mathbf{s}^\top$. For any $\mathbf{v} \in \mathbb{R}^n$,

$$\langle \mathbf{v}, (\mathrm{diag}(\mathbf{s}) - \mathbf{s}\mathbf{s}^\top)\mathbf{v}\rangle = \sum_i s_i v_i^2 - \left(\sum_i s_i v_i\right)^2,$$

which is the covariance over a probability vector $\mathbf{s}$, hence $\|\nabla \mathrm{softmax}\,(\mathbf{x})\,\|_2 \le \frac{1}{2}$. Thus,

$$\|\mathrm{softmax}\,(\mathbf{x}') - \mathrm{softmax}\,(\mathbf{x})\,\|_2 = \left\| \int_0^1 \nabla \mathrm{softmax}\,(\mathbf{x} + t(\mathbf{x}' - \mathbf{x}))\,(\mathbf{x}' - \mathbf{x}) \mathrm{d}t \right\|$$
$$\le \int_0^1 \|\nabla \mathrm{softmax}\,(\mathbf{x} + t(\mathbf{x}' - \mathbf{x}))\|\,\|\mathbf{x}' - \mathbf{x}\| \mathrm{d}t = \tfrac{1}{2}\|\mathbf{x} - \mathbf{x}'\|_2.$$

This completes the proof. ■

Lemma D.6 (Rademacher bound for $\mathscr{H}^{\mathrm{CES}} \cup \mathscr{H}^{\mathrm{Cobbs}} \cup \mathscr{H}^{\mathrm{Leon}}$)**.** *If Assumption D.1 holds, then*

$$\hat{\mathcal{R}}_K(\mathscr{H}^{\mathrm{CES}}) \le \sqrt{\frac{\pi n (D_{\mathbf{y}}^2 + n D_\sigma^2)(n + D_{\mathbf{p}}^2)}{8K}}.$$

The bound also holds for $\mathscr{H}^{\mathrm{Leon}}$ and $\mathscr{H}^{\mathrm{Cobbs}}$.

Proof. Observe that the expenditure share can be written as $\boldsymbol{\gamma}(\mathbf{p}) = \mathrm{softmax}\,(\boldsymbol{\Theta}\mathbf{q}(\mathbf{p}))$, where

$$\boldsymbol{\Theta} := [\mathrm{diag}(\mathbf{y}), \sigma \mathbf{I}_n] \in \mathbb{R}^{n \times 2n}, \quad \mathbf{q}(\mathbf{p}) := [\mathbf{1}; -\log(\mathbf{p})] \in \mathbb{R}^{2n}.$$

Under Assumption D.1, $|\sigma| \le D_\sigma$, so

$$\|\mathbf{\Theta}\|_F = \sqrt{\|\mathbf{y}\|_2^2 + n\sigma^2} \le \sqrt{D_{\mathbf{y}}^2 + nD_\sigma^2}, \quad \|\mathbf{q}(\mathbf{p})\|_2 = \sqrt{n + \|\log(\mathbf{p})\|_2^2} \le \sqrt{n + D_{\mathbf{p}}^2}.$$

By definition (29), $\boldsymbol{\xi}_1, \ldots, \boldsymbol{\xi}_K$ are i.i.d. Rademacher vectors. Hence, the concatenation $[\boldsymbol{\xi}_1; \ldots; \boldsymbol{\xi}_K]$ are i.i.d. Rademacher variables in dimension nK. Consider $\boldsymbol{\zeta}_k \sim \mathbf{N}(\mathbf{0}, \mathbf{I}_n)$ for all $k \in [K]$, and similarly $[\boldsymbol{\zeta}_1; \ldots; \boldsymbol{\zeta}_K]$ are i.i.d. Gaussian variables in dimension nK. For convenience, define

$$\mathbf{v} = [\operatorname{softmax}(\mathbf{\Theta}\mathbf{q}(\mathbf{p}_1)); \ldots; \operatorname{softmax}(\mathbf{\Theta}\mathbf{q}(\mathbf{p}_K))] \in \mathbb{R}^{nK},$$

so by Lemma D.3 applied to the class $\mathscr{H}^{\mathrm{CES}}$ at the points $\mathbf{p}_1, \ldots, \mathbf{p}_K$, we have the following inequality:

$$\begin{aligned} K \cdot \hat{\mathcal{R}}_K(\mathscr{H}^{\mathrm{CES}}) &= \mathbb{E}_{\boldsymbol{\xi}}\left[\sup_{\mathbf{\Theta}} \sum_{k=1}^K \langle \boldsymbol{\xi}_k, \operatorname{softmax}(\mathbf{\Theta}\mathbf{q}(\mathbf{p}_k))\rangle\right] \\ &\le \sqrt{\tfrac{\pi}{2}} \cdot \mathbb{E}_{\boldsymbol{\zeta}}\left[\sup_{\mathbf{\Theta}} \sum_{k=1}^K \langle \boldsymbol{\zeta}_k, \operatorname{softmax}(\mathbf{\Theta}\mathbf{q}(\mathbf{p}_k))\rangle\right]. \end{aligned} \tag{37}$$

To the fact that $\boldsymbol{\zeta}_1, ..., \boldsymbol{\zeta}_K$ are i.i.d. Gaussian variables, the following inequalities hold:

$$\mathbb{E}\left[\langle \boldsymbol{\zeta}_k, \boldsymbol{\zeta}_k\rangle\right] = \operatorname{tr}(\mathbb{E}[\boldsymbol{\zeta}_k \boldsymbol{\zeta}_k^\top]) = \operatorname{tr}(\mathbf{I}_n) = n \tag{38a}$$

$$\forall k \ne k', \quad \mathbb{E}\left[\langle \boldsymbol{\zeta}_k, \boldsymbol{\zeta}_{k'}\rangle\right] = 0, \tag{38b}$$

$$\forall \mathbf{v} \in \mathbb{R}^n, \quad \mathbb{E}\left(\langle \boldsymbol{\zeta}_k, \mathbf{v}\rangle^2\right) = \mathbb{E}\left(\operatorname{tr}\left(\boldsymbol{\zeta}_k \boldsymbol{\zeta}_k^\top \mathbf{v}\mathbf{v}^\top\right)\right) = \operatorname{tr}\left(\mathbb{E}\left(\boldsymbol{\zeta}_k \boldsymbol{\zeta}_k^\top\right) \mathbf{v}\mathbf{v}^\top\right) = \|\mathbf{v}\|_2^2. \tag{38c}$$

Define the Gaussian processes indexed by $\mathbf{\Theta}$,

$$X_{\mathbf{\Theta}} := \sum_{k=1}^K \langle \boldsymbol{\zeta}_k, \operatorname{softmax}(\mathbf{\Theta}\mathbf{q}(\mathbf{p}_k))\rangle, \quad Y_{\mathbf{\Theta}} := \tfrac{1}{2}\sum_{k=1}^K \langle \boldsymbol{\zeta}_k, \mathbf{\Theta}\mathbf{q}(\mathbf{p}_k)\rangle.$$

It is easy to see $\mathbb{E}[X_{\mathbf{\Theta}}] = \mathbb{E}[Y_{\mathbf{\Theta}}] = 0$. For any $\mathbf{\Theta}, \mathbf{\Theta}'$, let

$$\mathbf{v}_k := \operatorname{softmax}(\mathbf{\Theta}\mathbf{q}(\mathbf{p}_k)) - \operatorname{softmax}(\mathbf{\Theta}'\mathbf{q}(\mathbf{p}_k)).$$

So we have,

$$\begin{aligned} \mathbb{E}[(X_{\mathbf{\Theta}} - X_{\mathbf{\Theta}'})^2] &= \mathbb{E}\left[\sum_{k=1}^K \langle \boldsymbol{\zeta}_k, \operatorname{softmax}(\mathbf{\Theta}\mathbf{q}(\mathbf{p}_k))\rangle - \sum_{k=1}^K \langle \boldsymbol{\zeta}_k, \operatorname{softmax}\left(\mathbf{\Theta}'\mathbf{q}(\mathbf{p}_k)\right)\rangle\right]^2 \\ &= \mathbb{E}\left[\sum_{k=1}^K \langle \boldsymbol{\zeta}_k, \operatorname{softmax}(\mathbf{\Theta}\mathbf{q}(\mathbf{p}_k)) - \operatorname{softmax}\left(\mathbf{\Theta}'\mathbf{q}(\mathbf{p}_k)\right)\rangle\right]^2 \\ &= \mathbb{E}\left[\sum_{k=1}^K \langle \boldsymbol{\zeta}_k, \mathbf{v}_k\rangle\right]^2 = \mathbb{E}\left[\sum_{k,k'=1}^K \langle \boldsymbol{\zeta}_k, \mathbf{v}_k\rangle\langle \boldsymbol{\zeta}_{k'}, \mathbf{v}_{k'}\rangle\right] = \mathbb{E}\left[\sum_{k,k'=1}^K \operatorname{tr}\left(\boldsymbol{\zeta}_k \boldsymbol{\zeta}_{k'}^\top \mathbf{v}_{k'} \mathbf{v}_k^\top\right)\right] \\ &\underset{(38)}{=} \sum_{k=1}^K \mathbb{E}\langle \boldsymbol{\zeta}_k, \mathbf{v}_k\rangle^2 \underset{(38c)}{=} \sum_{k=1}^K \|\mathbf{v}_k\|_2^2 \underset{(36)}{\le} \tfrac{1}{4}\sum_{k=1}^K \|(\mathbf{\Theta} - \mathbf{\Theta}')\mathbf{q}(\mathbf{p}_k)\|_2^2 \\ &= \mathbb{E}\left[(Y_{\mathbf{\Theta}} - Y_{\mathbf{\Theta}'})^2\right]. \end{aligned}$$

By Sudakov-Fernique inequality; see, for example, Ledoux & Talagrand (1991, Theorem 3.15), we have

$$\mathbb{E}\left[\sup_{\mathbf{\Theta}} X_{\mathbf{\Theta}}\right] \le \mathbb{E}\left[\sup_{\mathbf{\Theta}} Y_{\mathbf{\Theta}}\right].$$

In view of (37), we have,

$$\hat{\mathcal{R}}_K(\mathscr{H}^{\mathrm{CES}}) \le \sqrt{\frac{\pi}{8}}\frac{1}{K}\mathbb{E}_{\boldsymbol{\zeta}}\left[\sup_{\mathbf{\Theta}} \sum_{k=1}^K \langle \boldsymbol{\zeta}_k, \mathbf{\Theta}\mathbf{q}(\mathbf{p}_k)\rangle\right].$$

The rest of the proof is to bound the RHS from Gaussian variables. Indeed, take $\mathbf{M} := \sum_{k=1}^{K} \mathbf{q}(\mathbf{p}_k)\boldsymbol{\zeta}_k^\top$,

$$\sup_{\boldsymbol{\Theta}} \sum_{k=1}^{K} \langle \boldsymbol{\zeta}_k, \boldsymbol{\Theta}\mathbf{q}(\mathbf{p}_k)\rangle = \sup_{\boldsymbol{\Theta}} \operatorname{tr}(\mathbf{M}\boldsymbol{\Theta}) \le \|\boldsymbol{\Theta}\|_F \|\mathbf{M}\|_F \le \sqrt{D_{\mathbf{y}}^2 + nD_\sigma^2}\|\mathbf{M}\|_F. \tag{39}$$

Note that,

$$\mathbb{E}\left[\|\mathbf{M}\|_F^2\right] = \mathbb{E}\left[\operatorname{tr}(\mathbf{M}^\top\mathbf{M})\right] = \sum_{k,k'=1}^{K} \langle \mathbf{q}(\mathbf{p}_k), \mathbf{q}(\mathbf{p}_{k'})\rangle \mathbb{E}\left[\langle \boldsymbol{\zeta}_k, \boldsymbol{\zeta}_{k'}\rangle\right]$$

$$\underset{(38)}{=} n\sum_{k=1}^{K} \|\mathbf{q}(\mathbf{p}_k)\|_2^2 \le nK(n + D_{\mathbf{p}}^2),$$

Combining, we have, $\hat{\mathcal{R}}_K(\mathscr{H}^{\mathrm{CES}}) \le \sqrt{\frac{\pi n(D_{\mathbf{y}}^2 + nD_\sigma^2)(n+D_{\mathbf{p}}^2)}{8K}}$. This completes the proof. ■

The bound has a polynomial envelope in $n, D_{\mathbf{y}}, D_\sigma, D_{\mathbf{p}}$ as used in (15).

D.2.2 Linear class

The linear class $\mathscr{H}^{\mathrm{Lin}}$ has finitely many distinct evaluations on any fixed sample, without requiring bounded parameters.

Lemma D.7 (Rademacher bound for $\mathscr{H}^{\mathrm{Lin}}$)**.** *For strictly positive sample prices and the smallest-index tie-breaking rule,*

$$\hat{\mathcal{R}}_K(\mathscr{H}^{\mathrm{Lin}}) \le \sqrt{\frac{2n\log\left(\frac{e(n-1)K}{2}\right)}{K}}. \tag{40}$$

Proof. By definition (29),

$$K \cdot \hat{\mathcal{R}}_K(\mathscr{H}^{\mathrm{Lin}}) = \mathbb{E}_{\boldsymbol{\xi}}\left[\sup_{\mathbf{y}\in\mathbb{R}^n} \sum_{k=1}^{K} \langle \boldsymbol{\xi}_k, \mathbf{e}_{j^*(\mathbf{y},\mathbf{p}_k)}\rangle\right]. \tag{41}$$

For fixed samples $\mathbf{p}_1, \ldots, \mathbf{p}_K$, introduce the mapping

$$\mathbf{J}(\mathbf{y}) := [j_1, \ldots, j_K] \in \mathcal{J} \subseteq [n]^K, \qquad j_k := j^*(\mathbf{y}, \mathbf{p}_k) = \min\arg\max_{j\in[n]}\{y_j - \log p_{k,j}\}.$$

The fixed smallest-index rule makes $\mathbf{J} : \mathbb{R}^n \to \mathcal{J}$ single-valued and piecewise constant; its value changes only when, for some sample k, two coordinates $j \ne j'$ become tied:

$$y_j - \log p_{k,j} = y_{j'} - \log p_{k,j'}.$$

That is, the values depends on the homogenous hyperplane in $\mathbf{y}$

$$H_{jj'k} := \left\{\mathbf{y} \in \mathbb{R}^n : y_j - y_{j'} = \log p_{k,j} - \log p_{k,j'}\right\}, \qquad j < j', k \in [K]. \tag{42}$$

In other words, by the affine hyperplanes (42), $\mathbb{R}^n$ is partitioned into finitely many regions, in each of which $\mathbf{J}$ is constant. This means the image $\mathcal{J}$ is finite. By Cover (1965) (see, also Matousek (2002, 6.1.1)), we can bound the number of possible regions generated by these hyperplanes, since there are at most $\binom{n}{2}K$ of them in all:

$$|\mathcal{J}| \le \sum_{i=0}^{n} \binom{\binom{n}{2}K}{i} \le \left(\frac{e\binom{n}{2}K}{n}\right)^n = \left(\frac{e\binom{n}{2}K}{n}\right)^n = \left(\frac{e(n-1)K}{2}\right)^n, \tag{43}$$

where the second inequality is due to Shalev-Shwartz & Ben-David (2014, Lemma A.5). Writing the Rademacher vectors into $\boldsymbol{\xi} := [\boldsymbol{\xi}_1; \cdots; \boldsymbol{\xi}_K] \in \{\pm 1\}^{nK}$, we can associate each labeling $\mathbf{J} \in \mathcal{J}$ with the vector

$$\mathbf{e}(\mathbf{J}) := [\mathbf{e}_{j_1}; \cdots; \mathbf{e}_{j_K}] \in \mathbb{R}^{nK}, \qquad \|\mathbf{e}(\mathbf{J})\|_2 = \sqrt{K}.$$

Hence, we have,

$$\mathbb{E}_{\boldsymbol{\xi}}\left[\sup_{\|\mathbf{y}\|_2\le D_{\mathbf{y}}}\sum_{k=1}^{K}\langle\boldsymbol{\xi}_k,\mathbf{e}_{j^*(\mathbf{y},\mathbf{p}_k)}\rangle\right]\underset{(\square_1)}{\le}\mathbb{E}_{\boldsymbol{\xi}}\left[\max_{\mathbf{J}\in\mathcal{J}}\langle\boldsymbol{\xi},\mathbf{e}(\mathbf{J})\rangle\right]\underset{(\square_2)}{\le}\sqrt{K}\cdot\sqrt{2\log|\mathcal{J}|}. \tag{44}$$

Here, part $\square_1$ is because, $\langle\boldsymbol{\xi}_k,\mathbf{e}_{j^*(\mathbf{y},\mathbf{p}_k)}\rangle$ depends on $\mathbf{y}$ only through $\mathbf{J}(\mathbf{y})$. Part $\square_2$ is due to the Massart lemma (Massart, 2000, Lemma 5.2). Substituting (43) into (44),

$$\hat{\mathcal{R}}_K(\mathscr{H}^{\mathrm{Lin}})\le\frac{\sqrt{K}\cdot\sqrt{2\log|\mathcal{J}|}}{K}=\sqrt{\frac{2\log|\mathcal{J}|}{K}}\le\sqrt{\frac{2n\log\left(\frac{e(n-1)K}{2}\right)}{K}}.$$

■

Clearly, it is bounded by a polynomial in n.

D.3 Generalization risk for constant wealth (WP_T)

We now combine vector contraction with the bound for a finite union of classes.

Theorem D.1 (Generalization bound under constant wealth)**.** *For any $\delta\in(0,1)$, with probability at least $1-\delta$ over the i.i.d. sampled prices $\mathbf{p}_1,\dots,\mathbf{p}_K\sim\tau$, the following bound holds simultaneously for every $\mathbf{h}\in\mathrm{conv}(\mathscr{H})$:*

$$\mathbb{E}_\tau[\ell(\mathbf{h}(\mathbf{p}),\mathbf{g}(\mathbf{p}))]\le L(\mathbf{h})+4\sqrt{\frac{\log\bar{f}}{K}}+2\sqrt{2}\max_{f\in[\bar{f}]}\hat{\mathcal{R}}_K(\mathscr{H}^f)+6\sqrt{\frac{\ln(2/\delta)}{2K}}. \tag{45}$$

Proof. By Theorem 3.3 in Mohri et al. (2018), for any $\delta\in(0,1)$, with probability at least $1-\delta$ over the sampled prices, the following bound holds simultaneously for every $\mathbf{h}\in\mathrm{conv}(\mathscr{H})$:

$$\mathbb{E}_\tau[\ell(\mathbf{h}(\mathbf{p}),\mathbf{g}(\mathbf{p}))]\le L(\mathbf{h})+2\hat{\mathcal{R}}_K(\ell\circ\mathrm{conv}(\mathscr{H}))+3\sup|\ell(\mathbf{h}(\mathbf{p}),\mathbf{g}(\mathbf{p}))|\sqrt{\frac{\ln(2/\delta)}{2K}}, \tag{46}$$

where the supremum is taken over $\mathbf{h}\in\mathrm{conv}(\mathscr{H})$ and $\mathbf{p}\in\mathcal{P}$, with the normalized market shares $\mathbf{g}(\mathbf{p})\in\Delta_n$. By Lemma D.2,

$$\hat{\mathcal{R}}_K(\ell\circ\mathrm{conv}(\mathscr{H}))\underset{(30)}{\le}\sqrt{2}\hat{\mathcal{R}}_K(\mathrm{conv}(\mathscr{H}))=\sqrt{2}\hat{\mathcal{R}}_K(\mathscr{H}), \tag{47}$$

Recall that (30) is by setting $\Phi_k(\cdot)=\ell(\cdot,\mathbf{g}_k)$, which is 1-Lipschitz in the first argument (cf. Lemma D.1). The second equation for the convex hull is well-known; see, e.g., Mohri et al. (2018, Lemma 7.4). Since $\mathbf{h}(\mathbf{p}),\mathbf{g}(\mathbf{p})\in\Delta_n$, it holds that

$$\begin{aligned}\sup|\ell|&=\sup\|\mathbf{h}(\mathbf{p})-\mathbf{g}(\mathbf{p})\|_\infty\\&\le\sup\|\mathbf{h}(\mathbf{p})-\mathbf{g}(\mathbf{p})\|_1\\&\le\sup\left(\|\mathbf{h}(\mathbf{p})\|_1+\|\mathbf{g}(\mathbf{p})\|_1\right)\\&=2.\end{aligned}$$

Combining,

$$\mathbb{E}_\tau[\ell(\mathbf{h}(\mathbf{p}),\mathbf{g}(\mathbf{p}))]\le L(\mathbf{h})+2\sqrt{2}\cdot\hat{\mathcal{R}}_K(\mathscr{H})+6\sqrt{\frac{\ln(2/\delta)}{2K}}. \tag{48}$$

By Lemma D.4,

$$\begin{aligned}2\sqrt{2}\hat{\mathcal{R}}_K(\mathscr{H})&\le2\sqrt{2}\left(\max_{f\in[\bar{f}]}\hat{\mathcal{R}}_K(\mathscr{H}^f)+\sqrt{\frac{2\log\bar{f}}{K}}\right)\\&=4\sqrt{\frac{\log\bar{f}}{K}}+2\sqrt{2}\max_{f\in[\bar{f}]}\hat{\mathcal{R}}_K(\mathscr{H}^f).\end{aligned}$$

Substituting into (48) proves (45). ■

D.4 GENERALIZATION RISK FOR THE EXTENSION ($\mathrm{WP}_{+,T}$)

We consider the joint wealth and supply formulation in Section 3.3, in which $\mathbf{a}$ is learned together with the android endowments. Fix a nonempty base class $\mathscr{H}$ of expenditure-share functions independently of the training sample. For normalized prices, the predictor is

$$\mathbf{h}(\mathbf{p}) = \sum_{t\in[T]} \langle \mathbf{p}, \mathbf{b}_t\rangle \boldsymbol{\gamma}_t(\mathbf{p}) - \mathrm{diag}(\mathbf{p})\mathbf{a},$$
$$\mathbf{a} \geq \mathbf{0}, \quad \mathbf{b}_t \geq \mathbf{0}, \quad \sum_{t\in[T]} \mathbf{b}_t = \mathbf{1} + \mathbf{a}, \quad \boldsymbol{\gamma}_t \in \mathscr{H}. \tag{49}$$

The vector $\mathbf{a}$ may vary with the training sample within this class. We first prove a bound uniform over the class, and then allow the supply bound itself to vary. Regrouping the androids by endowment coordinate yields

$$\hat{\boldsymbol{\gamma}}_j(\mathbf{p}) := \sum_{t\in[T]} \frac{b_{t,j}}{1+a_j} \boldsymbol{\gamma}_t(\mathbf{p}) \in \mathrm{conv}(\mathscr{H}),$$
$$\mathbf{h}(\mathbf{p}) = \sum_{j\in[n]} p_j \left[(1+a_j)\hat{\boldsymbol{\gamma}}_j(\mathbf{p}) - a_j \mathbf{e}_j\right]. \tag{50}$$

Lemma D.8 (Rademacher bound under linear wealth and extra supply)**.** *For every sample of K prices in Δ_n,*

$$\hat{\mathcal{R}}_K(\mathscr{Q}) \leq n(1+D_{\mathbf{a}})\hat{\mathcal{R}}_K(\mathscr{H}) + D_{\mathbf{a}}\sqrt{\frac{n}{2K}}. \tag{51}$$

In particular, at $D_{\mathbf{a}} = 0$,

$$\hat{\mathcal{R}}_K(\mathscr{Q}) \leq n\hat{\mathcal{R}}_K(\mathscr{H}). \tag{52}$$

Proof. Fix any $\boldsymbol{\gamma}_0 \in \mathscr{H}$ and define the random quantities

$$S_j := \sup_{\boldsymbol{\gamma}\in\mathrm{conv}(\mathscr{H})} \sum_{k=1}^{K} p_{k,j}\langle \boldsymbol{\xi}_k, \boldsymbol{\gamma}(\mathbf{p}_k) - \boldsymbol{\gamma}_0(\mathbf{p}_k)\rangle \geq 0,$$
$$Z_j := \sum_{k=1}^{K} p_{k,j}\langle \boldsymbol{\xi}_k, \boldsymbol{\gamma}_0(\mathbf{p}_k) - \mathbf{e}_j\rangle. \tag{53}$$

The choice $\boldsymbol{\gamma} = \boldsymbol{\gamma}_0$ ensures $S_j \geq 0$. An expected Rademacher supremum is convex and even in each sample coefficient, so its maximum over $[-1,1]^K$ is attained at a vertex, with the same value at every vertex. Since $p_{k,j} \in [0,1]$ and the fixed centering term has mean zero,

$$\begin{aligned}\mathbb{E}_{\boldsymbol{\xi}} S_j &= \mathbb{E}_{\boldsymbol{\xi}}\left[\sup_{\boldsymbol{\gamma}\in\mathrm{conv}(\mathscr{H})} \sum_{k=1}^{K} p_{k,j}\langle \boldsymbol{\xi}_k, \boldsymbol{\gamma}(\mathbf{p}_k)\rangle\right] \\ &\leq \mathbb{E}_{\boldsymbol{\xi}}\left[\sup_{\boldsymbol{\gamma}\in\mathrm{conv}(\mathscr{H})} \sum_{k=1}^{K} \langle \boldsymbol{\xi}_k, \boldsymbol{\gamma}(\mathbf{p}_k)\rangle\right] \\ &= K\hat{\mathcal{R}}_K(\mathscr{H}),\end{aligned} \tag{54}$$

where the final identity is the convex-hull identity (Mohri et al., 2018, Lemma 7.4). By (50),

$$\mathbf{h}(\mathbf{p}) = \boldsymbol{\gamma}_0(\mathbf{p}) + \sum_{j\in[n]} p_j \left[(1+a_j)(\hat{\boldsymbol{\gamma}}_j(\mathbf{p}) - \boldsymbol{\gamma}_0(\mathbf{p})) + a_j(\boldsymbol{\gamma}_0(\mathbf{p}) - \mathbf{e}_j)\right]. \tag{55}$$

Using $0 \leq a_j \leq D_{\mathbf{a}}$, $S_j \geq 0$, and the zero mean of the fixed $\boldsymbol{\gamma}_0$ term, we obtain

$$\begin{aligned}K\hat{\mathcal{R}}_K(\mathscr{Q}) &\leq (1+D_{\mathbf{a}}) \sum_{j\in[n]} \mathbb{E}_{\boldsymbol{\xi}} S_j + D_{\mathbf{a}} \sum_{j\in[n]} \mathbb{E}_{\boldsymbol{\xi}} \max\{Z_j, 0\} \\ &\leq n(1+D_{\mathbf{a}})K\hat{\mathcal{R}}_K(\mathscr{H}) + \frac{D_{\mathbf{a}}}{2} \sum_{j\in[n]} \mathbb{E}_{\boldsymbol{\xi}}|Z_j|,\end{aligned}$$

where the second inequality uses (54) and the symmetry of Z_j. Cauchy-Schwarz bounds the remaining term as

$$\begin{aligned}\sum_{j\in[n]} \mathbb{E}_{\boldsymbol{\xi}}|Z_j| &\le \sum_{j\in[n]} \sqrt{\sum_{k=1}^{K} p_{k,j}^2 \|\boldsymbol{\gamma}_0(\mathbf{p}_k) - \mathbf{e}_j\|_2^2} \\ &\underset{\square_1}{\le} \sqrt{2} \sum_{j\in[n]} \sqrt{\sum_{k=1}^{K} p_{k,j}^2} \\ &\underset{\square_2}{\le} \sqrt{2nK}.\end{aligned}$$

Here $\square_1$ uses $\|\boldsymbol{\gamma}_0(\mathbf{p}_k) - \mathbf{e}_j\|_2^2 \le 2$, and $\square_2$ is Cauchy-Schwarz with $\sum_k \|\mathbf{p}_k\|_2^2 \le K$. Dividing by K proves (51). ■

The additional complexity term accounts for learning $\mathbf{a}$.

Proof of Theorem 4.1 The nonnegative vectors $\sum_t \langle \mathbf{p}, \mathbf{b}_t\rangle \boldsymbol{\gamma}_t(\mathbf{p})$ and $\mathbf{g}(\mathbf{p}) + \mathrm{diag}(\mathbf{p})\mathbf{a}$ have the same total mass $1 + \langle \mathbf{p}, \mathbf{a}\rangle$. Consequently,

$$\begin{aligned}\ell(\mathbf{h}(\mathbf{p}), \mathbf{g}(\mathbf{p})) &= \left\| \sum_t \langle \mathbf{p}, \mathbf{b}_t\rangle \boldsymbol{\gamma}_t(\mathbf{p}) - \big(\mathbf{g}(\mathbf{p}) + \mathrm{diag}(\mathbf{p})\mathbf{a}\big) \right\|_\infty \\ &\le 1 + \langle \mathbf{p}, \mathbf{a}\rangle \\ &\le 1 + D_{\mathbf{a}}. \end{aligned} \tag{56}$$

Apply the Rademacher inequality used in (46) and the vector-contraction inequality (30) to $\mathscr{Q}$, and use Lemma D.8. By Lemma D.4,

$$\begin{aligned}2\sqrt{2}n(1 + D_{\mathbf{a}})\hat{\mathcal{R}}_K(\mathscr{H}) &\le 2\sqrt{2}n(1 + D_{\mathbf{a}}) \left(\max_{f\in[\bar{f}]} \hat{\mathcal{R}}_K(\mathscr{H}^f) + \sqrt{\frac{2\log \bar{f}}{K}} \right) \\ &= 2\sqrt{2}n(1 + D_{\mathbf{a}}) \max_{f\in[\bar{f}]} \hat{\mathcal{R}}_K(\mathscr{H}^f) + 4n(1 + D_{\mathbf{a}})\sqrt{\frac{\log \bar{f}}{K}}.\end{aligned}$$

Substituting this inequality and enlarging the confidence constant to the convention in (48) proves (15). ■

Under Assumption D.1, substituting the base-class estimates in Lemmas D.4, D.6 and D.7 gives a generalization gap of order $\widetilde{\mathcal{O}}((1 + D_{\mathbf{a}})n/\sqrt{K})$ for fixed dimension and parameter bounds, independently of the number of androids. The case $D_{\mathbf{a}} = 0$ recovers the linear-wealth bound without extra supply.

Remark 2 (A bound for the learned supply size)**.** *A simultaneous bound over increasing supply budgets allows* $\mathbf{a}$ *to be optimized without a prescribed cap. For each integer* $m \ge 0$*, set*

$$D_m := 2^m - 1, \qquad \alpha_m := \frac{\alpha}{(m+1)(m+2)}. \tag{57}$$

Apply Theorem 4.1 at each (D_m, α_m)*. Since* $\sum_{m\ge 0} \alpha_m = \alpha$*, these bounds hold simultaneously with probability at least* $1 - \alpha$*. For any learned finite supply vector, choose*

$$m(\mathbf{a}) := \lceil \log_2(1 + \|\mathbf{a}\|_\infty) \rceil . \tag{58}$$

Then (15) *applies with* $D_{\mathbf{a}} = D_{m(\mathbf{a})}$ *and confidence term*

$$6(1 + D_{m(\mathbf{a})}) \sqrt{\frac{\ln\big(2(m(\mathbf{a})+1)(m(\mathbf{a})+2)/\alpha\big)}{2K}}. \tag{59}$$

This bound holds uniformly for data-dependent $\mathbf{a}$*; its size grows with the learned extra supply.*

E CONVERGENCE ANALYSIS OF THE FRANK-WOLFE IMPLEMENTATION OF CRM

We prove Theorem 4.2 in three steps:

(1) We first reinterpret our decision variables using matrix notation (Section E.1), this enables us to simply the notations used in the convergence analysis.

(2) Next, we recall the standard convergence results of the Frank-Wolfe method for a generic smooth objective with an inexact linear minimization oracle (Section E.2, Proposition E.1). This is a standard result (Jaggi, 2013) and can be omitted for readers who are familiar with the method.

(3) Finally, we apply the results established in the previous two steps to show the complexity of CRM. Because our problem is nonsmooth, we apply Nesterov's smoothing technique here (Nesterov, 2005) (Section E.3). Then, we establish the connection of LMO used in the algorithm to SEP problem.

We then discuss how additive SEP errors affect the bound in Section E.4.

E.1 FEASIBLE PREDICTORS

We use one notation for the wealth specifications:

$$\begin{aligned} \mathcal{J} &:= [n], \quad a_{k,j} := p_{k,j} && \text{for linear wealth,} \\ \mathcal{J} &:= \{0\}, \quad a_{k,0} := 1 && \text{for constant wealth.} \end{aligned} \tag{60}$$

Here j indexes endowment coordinates, while expenditure shares remain n-dimensional. For an android $\boldsymbol{\gamma} \in \mathscr{H}$, define

$$\mathbf{A}_j(\boldsymbol{\gamma}) := [a_{1,j}\boldsymbol{\gamma}(\mathbf{p}_1), \dots, a_{K,j}\boldsymbol{\gamma}(\mathbf{p}_K)], \quad j \in \mathcal{J}, \tag{61}$$

$$\mathcal{D}_j := \operatorname{conv}\{\mathbf{A}_j(\boldsymbol{\gamma}) : \boldsymbol{\gamma} \in \mathscr{H}\}, \qquad \mathcal{D} := \sum_{j \in \mathcal{J}} \mathcal{D}_j, \tag{62}$$

where the sum is a Minkowski sum. For any finite collection of androids with coefficients $b_{s,j} \geq 0$ and $\sum_s b_{s,j} = 1$ for each $j \in \mathcal{J}$, the predictor satisfies

$$\mathbf{X} := [\mathbf{h}(\mathbf{p}_1), \dots, \mathbf{h}(\mathbf{p}_K)] = \sum_{j \in \mathcal{J}} \sum_s b_{s,j} \mathbf{A}_j(\boldsymbol{\gamma}_s) \in \mathcal{D}. \tag{63}$$

Conversely, pooling the androids in the finite convex combinations defining the sets $\mathcal{D}_j$ and using their coefficients realizes every $\mathbf{X} \in \mathcal{D}$. Since $\sum_{j \in \mathcal{J}} a_{k,j} = 1$, each column is a convex combination of expenditure shares, and hence

$$\mathbf{X}\mathbf{e}_k \in \Delta_n, \quad \mathbf{X} \in \mathcal{D},\ k \in [K], \qquad \sup_{\mathbf{X},\mathbf{Y} \in \mathcal{D}} \|\mathbf{X} - \mathbf{Y}\|_F^2 \leq 2K. \tag{64}$$

We next extend this representation to the joint linear-wealth and extra-supply formulation. Fix $D_{\mathbf{a}} \geq 0$, take $\mathcal{J} = [n]$, and define the fixed supplier matrices

$$\mathbf{P}_j := [p_{1,j}\mathbf{e}_j, \dots, p_{K,j}\mathbf{e}_j], \qquad j \in [n]. \tag{65}$$

For each endowment coordinate, define

$$\begin{aligned} \mathcal{D}_j^+(D_{\mathbf{a}}) &:= \{(1 + a_j)\mathbf{Y}_j - a_j\mathbf{P}_j : \mathbf{Y}_j \in \mathcal{D}_j,\ 0 \leq a_j \leq D_{\mathbf{a}}\} \\ &= \operatorname{conv}\left(\mathcal{D}_j \cup ((1 + D_{\mathbf{a}})\mathcal{D}_j - D_{\mathbf{a}}\mathbf{P}_j)\right), \\ \mathcal{D}^+(D_{\mathbf{a}}) &:= \sum_{j \in [n]} \mathcal{D}_j^+(D_{\mathbf{a}}). \end{aligned} \tag{66}$$

For $D_{\mathbf{a}} > 0$, the second equality follows by writing $a_j = \lambda D_{\mathbf{a}}$ and expressing the point as a convex combination of its two endpoint forms; conversely, convexity of $\mathcal{D}_j$ absorbs any convex combination of two endpoint forms into $(1 + a_j)\mathbf{Y}_j - a_j\mathbf{P}_j$. The equality is immediate when $D_{\mathbf{a}} = 0$. Indeed,

for any feasible endowments in ($\mathrm{WP}_{+,T}$), set

$$\mathbf{Y}_j := \frac{1}{1+a_j}\sum_t b_{t,j}\mathbf{A}_j(\boldsymbol{\gamma}_t) \in \mathcal{D}_j,$$
$$\mathbf{X} = [\mathbf{h}(\mathbf{p}_1),\dots,\mathbf{h}(\mathbf{p}_K)] = \sum_{j\in[n]}((1+a_j)\mathbf{Y}_j - a_j\mathbf{P}_j) \in \mathcal{D}^+(D_{\mathbf{a}}). \tag{67}$$

Conversely, pooling the finite convex combinations defining the $\mathbf{Y}_j$ realizes every point of $\mathcal{D}^+(D_{\mathbf{a}})$ by android endowments and some $\mathbf{0} \le \mathbf{a} \le D_{\mathbf{a}}\mathbf{1}$. For a column $\mathbf{x} = \mathbf{X}\mathbf{e}_k$ of a feasible predictor, write $\mathbf{x} = \mathbf{u} - \mathbf{v}$, where $\mathbf{u} \ge \mathbf{0}$ has total mass $1 + \langle \mathbf{p}_k, \mathbf{a}\rangle$ and $\mathbf{v} = \mathrm{diag}(\mathbf{p}_k)\mathbf{a}$ has total mass $\langle \mathbf{p}_k, \mathbf{a}\rangle \le D_{\mathbf{a}}$. Thus $\langle \mathbf{1}, \mathbf{x}\rangle = 1$ and $\|\mathbf{x}\|_1 \le 1 + 2D_{\mathbf{a}}$. For any two feasible columns, their difference has zero sum, so its squared Euclidean norm is at most one half of its squared ℓ_1 norm. Consequently,

$$\sup_{\mathbf{X},\mathbf{Y}\in\mathcal{D}^+(D_{\mathbf{a}})} \|\mathbf{X}-\mathbf{Y}\|_F^2 \le 2K(1+2D_{\mathbf{a}})^2. \tag{68}$$

The empirical loss is

$$L(\mathbf{X}) := \frac{1}{K}\sum_{k\in[K]} \|\mathbf{X}\mathbf{e}_k - \mathbf{g}_k\|. \tag{69}$$

Define its infimum over the bounded extra-supply domain by

$$L_+^* := \inf_{\mathbf{X}\in\mathcal{D}^+(D_{\mathbf{a}})} L(\mathbf{X}). \tag{70}$$

At $D_{\mathbf{a}} = 0$, its infimum over $\mathcal{D}$ is L_+^* for linear wealth and L^* for constant wealth. The linear-wealth feasible set contains the constant-wealth one, so $L_+^* \le L^*$.

E.2 Frank-Wolfe with an inexact linear minimization oracle

Let $\mathcal{C}$ be a nonempty bounded convex set in a matrix space, and let F be convex with a β-Lipschitz gradient on $\mathcal{C}$ in the Frobenius norm. Define

$$F^* := \inf_{\mathbf{X}\in\mathcal{C}} F(\mathbf{X}), \qquad R := \sup_{\mathbf{X},\mathbf{Y}\in\mathcal{C}} \|\mathbf{X}-\mathbf{Y}\|_F. \tag{71}$$

At iteration t, compute the gradient

$$\mathbf{V}_t := \nabla F(\mathbf{X}_{t-1}). \tag{72}$$

An inexact linear minimization oracle (LMO) approximately minimizes $\langle \mathbf{V}_t, \mathbf{S}\rangle$ over $\mathbf{S} \in \mathcal{C}$, returning $\mathbf{S}_t \in \mathcal{C}$. Its additive error is

$$\varepsilon_t := \langle \mathbf{V}_t, \mathbf{S}_t\rangle - \inf_{\mathbf{S}\in\mathcal{C}} \langle \mathbf{V}_t, \mathbf{S}\rangle \ge 0, \tag{73}$$

where $\langle\cdot,\cdot\rangle$ is the Frobenius inner product. The standard Frank-Wolfe update is

$$\eta_t := \frac{2}{t+1}, \qquad \mathbf{X}_t := (1-\eta_t)\mathbf{X}_{t-1} + \eta_t\mathbf{S}_t. \tag{74}$$

Algorithm 2: Frank-Wolfe with an inexact LMO

Input: Smooth objective F; initial point $\mathbf{X}_0 \in \mathcal{C}$; an inexact LMO over $\mathcal{C}$; iteration budget T.
for $t = 1,\dots,T$ **do**
  Set $\mathbf{V}_t := \nabla F(\mathbf{X}_{t-1})$;
  Call the LMO to obtain $\mathbf{S}_t \in \mathcal{C}$;
  Update $\mathbf{X}_t$ using (74);
end
Output: $\mathbf{X}_T$.

Proposition E.1 (Inexact Frank-Wolfe)**.** *Under the stated assumptions, Algorithm 2 satisfies*

$$0 \le F(\mathbf{X}_T) - F^* \le \frac{2\beta R^2}{T+2} + E_T, \qquad T \ge 1, \tag{75}$$

where the weighted average of the oracle errors is

$$E_T := \frac{2}{T(T+1)} \sum_{t=1}^{T} t\varepsilon_t. \tag{76}$$

If $\varepsilon_t \le \delta$ at every iteration, then $E_T \le \delta$.

Proof. The curvature of F is at most βR^2. The descent estimate of Jaggi (2013, Lemma 5) implies

$$h_t := F(\mathbf{X}_t) - F^*,$$
$$h_t \le (1-\eta_t)h_{t-1} + \eta_t\varepsilon_t + \frac{\beta R^2}{2}\eta_t^2. \tag{77}$$

Unrolling the recurrence with (74) yields (75), with weight $\frac{2t}{T(T+1)}$ on ε_t. These weights sum to one, proving the last claim. ■

E.3 Application to CRM

Now, let us apply the results established in the previous two steps to show the complexity of CRM. Because our problem is nonsmooth, we apply Nesterov's smoothing technique here (Nesterov, 2005) (Section E.3).

Smoothed objective. For the nonzero seminorm in (69), define

$$\mathcal{Q} := \{\mathbf{U} \in \mathbb{R}^{n\times K} : \|\mathbf{U}\mathbf{e}_k\|_* \le 1,\ k \in [K]\}. \tag{78}$$

By the dual representation of the seminorm,

$$\begin{aligned} L(\mathbf{X}) &= \frac{1}{K}\sum_{k\in[K]} \|\mathbf{g}_k - \mathbf{X}\mathbf{e}_k\| \\ &= \frac{1}{K}\sum_{k\in[K]} \max_{\|\mathbf{u}_k\|_*\le 1} \langle \mathbf{u}_k, \mathbf{g}_k - \mathbf{X}\mathbf{e}_k\rangle \\ &= \frac{1}{K}\max_{\mathbf{U}\in\mathcal{Q}} \sum_{k\in[K]} \langle \mathbf{U}\mathbf{e}_k, \mathbf{g}_k - \mathbf{X}\mathbf{e}_k\rangle, \end{aligned} \tag{79}$$

where the last equality uses the independent constraints on the columns $\mathbf{U}\mathbf{e}_k = \mathbf{u}_k$. We instead apply Algorithm 2 with $\mathcal{C} \leftarrow \mathcal{D}^+(D_\mathbf{a})$ and $F \leftarrow L_\nu$, using the smoothing technique (Nesterov, 2005):

$$\begin{aligned} L_\nu(\mathbf{X}) &:= \frac{1}{K}\max_{\mathbf{U}\in\mathcal{Q}} \left\{ \sum_{k\in[K]} \langle \mathbf{U}\mathbf{e}_k, \mathbf{g}_k - \mathbf{X}\mathbf{e}_k\rangle - \frac{\nu}{2}\|\mathbf{U}\|_F^2 \right\}, \quad \nu > 0, \\ \varsigma_* &:= \max_{\|\mathbf{u}\|_*\le 1} \|\mathbf{u}\|_2. \end{aligned} \tag{80}$$

By Nesterov (2005, Theorem 1 and Eq. (2.7)), L_ν is $\frac{1}{\nu K}$-smooth and satisfies

$$L_\nu(\mathbf{X}) \le L(\mathbf{X}) \le L_\nu(\mathbf{X}) + \frac{\nu\varsigma_*^2}{2}. \tag{81}$$

The gradient of L_ν is computed from the unique maximizer as

$$\begin{aligned} \mathbf{U}_\nu(\mathbf{X})\mathbf{e}_k &:= \mathrm{proj}_{\{\mathbf{u}:\|\mathbf{u}\|_*\le 1\}}\left(\frac{\mathbf{g}_k - \mathbf{X}\mathbf{e}_k}{\nu}\right), \quad k \in [K], \\ \nabla L_\nu(\mathbf{X}) &= -\frac{1}{K}\mathbf{U}_\nu(\mathbf{X}), \end{aligned} \tag{82}$$

where proj denotes Euclidean projection. For the ∞-norm used in the main text, $\varsigma_* = 1$.

Linear minimization via SEP. Set the separation weights and define the global SEP values by

$$\begin{aligned} \mathbf{U}_t &:= -K\mathbf{V}_t = \mathbf{U}_\nu(\mathbf{X}_{t-1}), \\ \pi_j(\mathbf{U}) &:= \sup_{\boldsymbol{\gamma}\in\mathscr{H}} \sum_{k\in[K]} a_{k,j}\langle \mathbf{U}\mathbf{e}_k, \boldsymbol{\gamma}(\mathbf{p}_k)\rangle, \quad j \in \mathcal{J}. \end{aligned} \tag{83}$$

By (62), (61), and (83),

$$\begin{aligned}\inf_{\mathbf{S}\in\mathcal{D}}\langle\mathbf{V}_t,\mathbf{S}\rangle &= \inf_{\substack{\mathbf{S}_j\in\mathcal{D}_j\\ j\in\mathcal{J}}}\sum_{j\in\mathcal{J}}\langle\mathbf{V}_t,\mathbf{S}_j\rangle\\ &= \sum_{j\in\mathcal{J}}\inf_{\mathbf{S}_j\in\mathcal{D}_j}\langle\mathbf{V}_t,\mathbf{S}_j\rangle\\ &= \sum_{j\in\mathcal{J}}\inf_{\boldsymbol{\gamma}\in\mathscr{H}}\langle\mathbf{V}_t,\mathbf{A}_j(\boldsymbol{\gamma})\rangle\\ &= \sum_{j\in\mathcal{J}}\inf_{\boldsymbol{\gamma}\in\mathscr{H}}\sum_{k\in[K]}a_{k,j}\langle\mathbf{V}_t\mathbf{e}_k,\boldsymbol{\gamma}(\mathbf{p}_k)\rangle\\ &= -\frac{1}{K}\sum_{j\in\mathcal{J}}\sup_{\boldsymbol{\gamma}\in\mathscr{H}}\sum_{k\in[K]}a_{k,j}\langle\mathbf{U}_t\mathbf{e}_k,\boldsymbol{\gamma}(\mathbf{p}_k)\rangle\\ &= -\frac{1}{K}\sum_{j\in\mathcal{J}}\pi_j(\mathbf{U}_t). \end{aligned}\tag{84}$$

The choices of $\mathbf{S}_j\in\mathcal{D}_j$ are independent, and a linear functional has the same infimum over a convex hull as over its generating set. Thus the LMO without extra supply is equivalent to independent SEP calls indexed by $\mathcal{J}$.

For the bounded extra-supply domain, define

$$\begin{aligned}c_j(\mathbf{U}) &:= \sum_{k\in[K]}p_{k,j}\langle\mathbf{U}\mathbf{e}_k,\mathbf{e}_j\rangle,\\ \Pi_{j,D_\mathbf{a}}(\mathbf{U}) &:= \max\left\{\pi_j(\mathbf{U}),(1+D_\mathbf{a})\pi_j(\mathbf{U})-D_\mathbf{a}c_j(\mathbf{U})\right\}.\end{aligned}\tag{85}$$

By the endpoint representation in (66),

$$\inf_{\mathbf{S}\in\mathcal{D}^+(D_\mathbf{a})}\langle\mathbf{V}_t,\mathbf{S}\rangle = -\frac{1}{K}\sum_{j\in[n]}\Pi_{j,D_\mathbf{a}}(\mathbf{U}_t).\tag{86}$$

Hence the joint optimization of $\mathbf{a}$ requires the same SEP value $\pi_j(\mathbf{U}_t)$ as the no-extra-supply problem, followed by a comparison of the two endpoints $a_j\in\{0,D_\mathbf{a}\}$. Given an inexact SEP solution $\boldsymbol{\gamma}_{t,j}$, define its feasible score

$$\begin{aligned}\hat{\pi}_{t,j} &:= \sum_{k\in[K]}p_{k,j}\langle\mathbf{U}_t\mathbf{e}_k,\boldsymbol{\gamma}_{t,j}(\mathbf{p}_k)\rangle,\\ d_{t,j} &\in \operatorname{argmax}_{d\in\{0,D_\mathbf{a}\}}\left\{(1+d)\hat{\pi}_{t,j}-dc_j(\mathbf{U}_t)\right\},\\ \mathbf{S}_t &:= \sum_{j\in[n]}\left((1+d_{t,j})\mathbf{A}_j(\boldsymbol{\gamma}_{t,j})-d_{t,j}\mathbf{P}_j\right).\end{aligned}\tag{87}$$

Initialize Algorithm 2 at any feasible predictor with $\mathbf{a}=\mathbf{0}$. The update (74) preserves $\mathbf{X}_t\in\mathcal{D}^+(D_\mathbf{a})$ and jointly updates the android endowments and extra supply. After T iterations, the output is $\mathbf{X}_T$ with empirical risk

$$L_{+,T} := L(\mathbf{X}_T).\tag{88}$$

Accumulated oracle error. By (86), the LMO errors in (73) satisfy

$$\begin{aligned}\varepsilon_t &:= \frac{1}{K}\sum_{j\in[n]}\left[\Pi_{j,D_\mathbf{a}}(\mathbf{U}_t)-\left((1+d_{t,j})\hat{\pi}_{t,j}-d_{t,j}c_j(\mathbf{U}_t)\right)\right],\\ E_T &:= \frac{2}{T(T+1)}\sum_{t=1}^{T}t\varepsilon_t.\end{aligned}\tag{89}$$

Every difference is measured relative to the global SEP optimum over $\mathscr{H}$.

Choose the smoothing level fixed throughout the run as

$$\nu_{D_{\mathbf{a}}} := \frac{2\sqrt{2}(1+2D_{\mathbf{a}})}{\varsigma_*\sqrt{T+2}}. \tag{90}$$

The no-extra-supply specialization is

$$\nu_0 := \frac{2\sqrt{2}}{\varsigma_*\sqrt{T+2}}. \tag{91}$$

Proof of Theorem 4.2 For $L_{\nu_{D_{\mathbf{a}}}}$ on $\mathcal{D}^+(D_{\mathbf{a}})$, the smoothness and diameter bounds are $\beta = \frac{1}{\nu_{D_{\mathbf{a}}}K}$ and $R^2 \le 2K(1+2D_{\mathbf{a}})^2$. By (88), (81), and Proposition E.1,

$$\begin{aligned} L_{+,T} - L_+^* &\le L_{\nu_{D_{\mathbf{a}}}}(\mathbf{X}_T) - \inf_{\mathbf{X}\in\mathcal{D}^+(D_{\mathbf{a}})} L_{\nu_{D_{\mathbf{a}}}}(\mathbf{X}) + \frac{\nu_{D_{\mathbf{a}}}\varsigma_*^2}{2} \\ &\le \frac{4(1+2D_{\mathbf{a}})^2}{\nu_{D_{\mathbf{a}}}(T+2)} + E_T + \frac{\nu_{D_{\mathbf{a}}}\varsigma_*^2}{2}. \end{aligned} \tag{92}$$

Substitution yields

$$0 \le L_{+,T} - L_+^* \le \frac{2\sqrt{2}\varsigma_*(1+2D_{\mathbf{a}})}{\sqrt{T+2}} + E_T. \tag{93}$$

For the ∞-norm, $\varsigma_* = 1$, which proves (17). ■

Remark 3 (Number of androids)**.** *Each iteration of the CRM realization of Algorithm 2 makes n SEP calls in (87), adding at most n new androids. Starting from one initial android, the collected set therefore contains at most $1+nT$ androids after T iterations. Taking $T \ge \lceil 8\varsigma_*^2(1+2D_{\mathbf{a}})^2/\epsilon^2\rceil$ and the smoothing level in (90), (93) ensures excess risk at most $\epsilon + E_T$. For the ∞-norm loss, this requires $\mathcal{O}(n(1+2D_{\mathbf{a}})^2\epsilon^{-2})$ androids; with exact SEP, the excess risk is at most ϵ.*

E.4 Effect of additive SEP errors

Suppose every SEP call has additive error at most $\delta_{\mathrm{SEP}} \ge 0$ relative to its global optimum. Because each endpoint score in (85) has slope at most $1+D_{\mathbf{a}}$ in $\pi_j(\mathbf{U}_t)$, each bracketed gap in (89) is at most $(1+D_{\mathbf{a}})\delta_{\mathrm{SEP}}$. There are n calls per Frank-Wolfe iteration, and E_T is a weighted average of the iteration errors, so

$$\begin{aligned} \varepsilon_t &\le \frac{n(1+D_{\mathbf{a}})\delta_{\mathrm{SEP}}}{K}, \quad t\in[T], \\ E_T &\le \max_{t\in[T]}\varepsilon_t \le \frac{n(1+D_{\mathbf{a}})\delta_{\mathrm{SEP}}}{K}. \end{aligned} \tag{94}$$

Substituting (94) into (93), $T \ge \lceil 8\varsigma_*^2(1+2D_{\mathbf{a}})^2/\epsilon^2\rceil$ ensures excess risk at most $\epsilon + \frac{n(1+D_{\mathbf{a}})\delta_{\mathrm{SEP}}}{K}$.

F Proof of results on misspecification error I

We prove Theorem F.3, the formal version of Theorem 4.3, in three steps:

(1) We first assume that the excess demand of the real market is homogeneous of degree zero and satisfies Walras's law (Assumption F.1), so that prices can be restricted to Δ_n.

(2) We then show that this assumption is without loss of generality: any real market with arbitrary continuous wealth functions can be lifted to a market with one more good and one more agent that satisfies Assumption F.1 and preserves the Walrasian equilibria (Section F.1, Theorem F.1).

(3) Finally, if the real aggregate demand $\mathbf{d}$ is twice continuously differentiable, so is the lifted excess demand $\bar{\mathbf{z}}$. We apply Mantel's theorem (Theorem F.2) to $\bar{\mathbf{z}}$ on a compact set $\mathcal{K} \subset \mathrm{int}(\bar{\mathcal{P}})$ containing the sampled prices, which yields an exchange economy of $n+1$ homothetic consumers. Its wealth functions are linear, and the supplier correction yields zero misspecification error (Section F.2).

We begin by introducing the following assumption on the real market.

Assumption F.1. *The excess demand (cf. Equation* (4)*) in the real market is homogeneous of degree zero:* $\mathbf{z}(\lambda\mathbf{p}) = \mathbf{z}(\mathbf{p})$ *for all* $\lambda > 0$ *and* $\mathbf{p} \in \mathcal{P}$*, and it satisfies the Walras's law:* $\langle \mathbf{p}, \mathbf{z}(\mathbf{p})\rangle = 0$ *for all* $\mathbf{p} \in \mathcal{P}$*.*

The above assumptions can help limit the discussion to $\mathcal{P} = \Delta_n$ instead of a convex set; consequently, the total wealth can also be normalized to 1:

$$0 = \langle \mathbf{p}, \mathbf{z}(\mathbf{p})\rangle = \langle \mathbf{p}, \sum_{i\in[m]} \mathbf{x}_i(\mathbf{p}) - \mathbf{1}\rangle = \sum_{i\in[m]} w_i(\mathbf{p}) - 1.$$

However, this assumption holds in the Arrow-Debreu economy, but fails at constant and nonlinear wealth functions. For example, if all $\{w_i\}_{i\in[m]}$ are constant, it is clear that the demand at $\lambda\mathbf{p}, \lambda > 1$ will be smaller than $\mathbf{d}(\mathbf{p})$ because everything becomes more expensive. Fortunately, we can transform any real market to a lifted market that introduces money as the numéraire together with an additional agent, thereby satisfying the required conditions.

F.1 Transforming real market to lifted market

We describe the construction of the lifted market first, then, a formal version is provided. The price space $\mathcal{P} \subseteq \mathbb{R}^n_+$ of the original market is a compact convex set, and a price $\mathbf{q} \in \mathcal{P}$ may be taken at any scale. Its continuous aggregate demand $\mathbf{d} : \mathcal{P} \mapsto \mathbb{R}^n_+$ need not be homogeneous of degree zero, since the wealth functions $\{w_i\}$ are arbitrary; that is, $\mathbf{d}(\lambda\mathbf{q})$ and $\mathbf{d}(\mathbf{q})$ may differ. We use the numéraire trick: take money as the $(n+1)$-th commodity and define the lifted market as follows.

Money supply. By local non-satiation each agent exhausts its budget, so the value of aggregate demand equals the total wealth, $\langle \mathbf{q}, \mathbf{d}(\mathbf{q})\rangle = \sum_i w_i(\mathbf{q}) = W(\mathbf{q})$. We fix a money supply that dominates the wealth in excess of the value of supply,

$$M \geq \max\left\{0,\ \sup_{\mathbf{q}\in\mathcal{P}} \left(W(\mathbf{q}) - \langle \mathbf{q}, \mathbf{1}\rangle\right)\right\} \geq 0. \tag{95}$$

Since $\mathcal{P}$ is compact and W is continuous, the supremum in (95) is finite, so M can be chosen finite.

Lifted prices and demand. For $\mathbf{q} \in \mathcal{P}$ define the normalized lifted price and the lifted price set together with its positive cone

$$\bar{\mathbf{p}}(\mathbf{q}) := \frac{[\mathbf{q}; 1]}{1 + \langle \mathbf{q}, \mathbf{1}\rangle} \in \bar{\mathcal{P}}, \qquad \bar{\mathcal{P}} \qquad := \{\bar{\mathbf{p}}(\mathbf{q}) : \mathbf{q} \in \mathcal{P}\} \subseteq \Delta_{n+1},$$
$$\bar{\mathcal{C}} := \{\lambda\bar{\mathbf{p}}(\mathbf{q}) : \mathbf{q} \in \mathcal{P},\ \lambda > 0\}.$$

Writing $\bar{\mathbf{p}} = [\mathbf{p}; \pi]$, the money price is $\pi = \frac{1}{1+\langle \mathbf{q}, \mathbf{1}\rangle} > 0$ and $\mathbf{q} = \frac{\mathbf{p}}{\pi}$ recovers the original price. We define the lifted demand on $\bar{\mathcal{C}}$ through this ratio:

$$\bar{\mathbf{d}}(\bar{\mathbf{p}}) := [\mathbf{d}(\mathbf{q});\ d_{n+1}(\mathbf{q})],$$
$$d_{n+1}(\mathbf{q}) := M + \langle \mathbf{q}, \mathbf{1}\rangle - W(\mathbf{q}). \tag{96}$$

By (95), $d_{n+1}(\mathbf{q}) \geq 0$, so $\bar{\mathbf{d}}(\bar{\mathbf{p}}) \in \mathbb{R}^{n+1}_+$.

Lifted agents. To ensure $\bar{\mathbf{d}}$ can be rationalized by rational agents, keep each original agent $i \in [m]$ with the same utility and a rescaled budget, and add a single *rentier*, an agent indexed 0 who only cares about money:

$$\bar{u}_i(\mathbf{x}, m) := u_i(\mathbf{x}), \quad \bar{w}_i(\bar{\mathbf{p}}) := \pi\, w_i(\mathbf{q}),$$
$$\bar{u}_0(\mathbf{x}, m) := m, \quad \bar{w}_0(\bar{\mathbf{p}}) := \pi\, d_{n+1}(\mathbf{q}).$$

Each $\bar{u}_i$ inherits local non-satiation and concavity from u_i, and $\bar{u}_0(\mathbf{x}, m) = m$ is linear, hence locally non-satiated and concave. It is easy to see that any agent $i \in [m]$ has no interest in money:

$$\bar{\mathbf{x}}_i = \arg\max_{\bar{\mathbf{x}}\in\mathbb{R}^{n+1}_+} \bar{u}_i(\bar{\mathbf{x}}, m) \ \text{ s.t. } \ \langle \bar{\mathbf{p}}, \bar{\mathbf{x}}\rangle \leq \bar{w}_i(\bar{\mathbf{p}})$$
$$= \arg\max_{\bar{\mathbf{x}}\in\mathbb{R}^{n+1}_+} u_i(\mathbf{x}) \ \text{ s.t. } \ \langle \mathbf{p}, \mathbf{x}\rangle + \pi m \leq \pi w_i(\mathbf{q}) = [\mathbf{x}_i(\mathbf{q}); 0];$$

For the rentier,

$$\begin{aligned}\bar{\mathbf{x}}_0 &= \arg\max_{\bar{\mathbf{x}}\in\mathbb{R}^{n+1}_+} \bar{u}_0(\bar{\mathbf{x}}, m) \ \text{ s.t. } \ \langle \bar{\mathbf{p}}, \bar{\mathbf{x}}\rangle \le \bar{w}_0(\bar{\mathbf{p}}) \\ &= \arg\max_{\bar{\mathbf{x}}\in\mathbb{R}^{n+1}_+} m \ \text{ s.t. } \ \langle \mathbf{p}, \mathbf{x}\rangle + \pi m \le \pi d_{n+1}(\mathbf{q}) = [\mathbf{0}; d_{n+1}(\mathbf{q})].\end{aligned}$$

Therefore, the aggregate demand of the lifted agents is $\bar{\mathbf{d}}$.

Theorem F.1. *Let the real market have arbitrary continuous wealth functions $\{w_i\}_{i\in[m]}$, with aggregate demand $\mathbf{d}$ on the compact set $\mathcal{P} \subseteq \mathbb{R}^n_+$. The lifted market (96) satisfies*

(1) *The normalized prices lie in $\bar{\mathcal{P}} \subseteq \Delta_{n+1}$, and $\bar{\mathbf{d}} : \bar{\mathcal{C}} \mapsto \mathbb{R}^{n+1}_+$. Let $\bar{\mathbf{z}} : \bar{\mathcal{C}} \mapsto \mathbb{R}^{n+1}$ be the excess demand function:*

$$\bar{\mathbf{z}}(\bar{\mathbf{p}}) = \bar{\mathbf{d}}(\bar{\mathbf{p}}) - [\mathbf{1}; M].$$

Then, it is homogeneous of degree zero and satisfies the Walras's law.

(2) *If $[\mathbf{p}^*; \pi^*] \in \operatorname{int}(\bar{\mathcal{P}})$ is a Walrasian equilibrium of the lifted market, then $\mathbf{q}^* = \frac{\mathbf{p}^*}{\pi^*}$ is a Walrasian equilibrium of the original market. The reverse direction is also true.*

(3) *If $\mathcal{P} \subset \mathbb{R}^n_{++}$ and $\mathbf{d}$ is twice continuously differentiable on an open neighborhood of $\mathcal{P}$, then $\bar{\mathbf{z}}|_{\bar{\mathcal{P}}}$ admits a twice continuously differentiable extension to Δ_{n+1} that satisfies Walras's law.*

Proof. Nonnegativity of $\bar{\mathbf{d}}$ was shown above. For homogeneity, the map $\bar{\mathbf{p}} = [\mathbf{p}; \pi] \mapsto \mathbf{q} = \mathbf{p}/\pi$ is invariant under $\bar{\mathbf{p}} \mapsto \lambda\bar{\mathbf{p}}$ for $\lambda > 0$, so $\bar{\mathbf{d}}(\lambda\bar{\mathbf{p}}) = \bar{\mathbf{d}}(\bar{\mathbf{p}})$ and

$$\bar{\mathbf{z}}(\lambda\bar{\mathbf{p}}) = \bar{\mathbf{d}}(\lambda\bar{\mathbf{p}}) - [\mathbf{1}; M] = \bar{\mathbf{z}}(\bar{\mathbf{p}}). \tag{97}$$

For Walras's law, write $\bar{\mathbf{p}} = [\mathbf{p}; \pi]$ with $\mathbf{p} = \pi\mathbf{q}$, so that

$$\begin{aligned}\tfrac{1}{\pi}\langle \bar{\mathbf{p}}, \bar{\mathbf{z}}(\bar{\mathbf{p}})\rangle &= \langle [\mathbf{q}; 1], [\mathbf{d}(\mathbf{q}); d_{n+1}(\mathbf{q})] - [\mathbf{1}; M]\rangle \\ &= \langle [\mathbf{q}; 1], [\mathbf{d}(\mathbf{q}); d_{n+1}(\mathbf{q})]\rangle - \langle [\mathbf{q}; 1], [\mathbf{1}; M]\rangle \\ &= \langle \mathbf{q}, \mathbf{d}(\mathbf{q})\rangle + d_{n+1}(\mathbf{q}) - \langle \mathbf{q}, \mathbf{1}\rangle - M \\ &\underset{(96)}{=} \langle \mathbf{q}, \mathbf{d}(\mathbf{q})\rangle + (M + \langle \mathbf{q}, \mathbf{1}\rangle - \langle \mathbf{q}, \mathbf{d}(\mathbf{q})\rangle) - \langle \mathbf{q}, \mathbf{1}\rangle - M = 0,\end{aligned} \tag{98}$$

and $\pi > 0$ gives $\langle \bar{\mathbf{p}}, \bar{\mathbf{z}}(\bar{\mathbf{p}})\rangle = 0$. This proves part (1).

For part (2), the first n coordinates of $\bar{\mathbf{d}}(\bar{\mathbf{p}}^*) = [\mathbf{1}; M]$ read $\mathbf{d}(\mathbf{q}^*) = \mathbf{1}$; and when $\mathbf{d}(\mathbf{q}^*) = \mathbf{1}$, the last coordinate satisfies $d_{n+1}(\mathbf{q}^*) = M + \langle \mathbf{q}^*, \mathbf{1}\rangle - \langle \mathbf{q}^*, \mathbf{1}\rangle = M$ automatically, so the money market clears as well. The two clearing conditions are therefore equivalent. Since $\pi^* > 0$, the ray $\mathbf{q}^* = \frac{\mathbf{p}^*}{\pi^*}$ is well defined and lies in $\mathcal{P}$. Clearly, if $\mathbf{q}^* \in \mathcal{P}$ is an equilibrium, then $\bar{\mathbf{p}}^* = [\frac{\mathbf{q}^*}{1+\langle \mathbf{q}^*, \mathbf{1}\rangle}; \frac{1}{1+\langle \mathbf{q}^*, \mathbf{1}\rangle}] \in \operatorname{int}(\bar{\mathcal{P}})$ is an equilibrium of the lifted market.

For part (3), let $\mathcal{U} \subset \mathbb{R}^n_{++}$ be an open neighborhood of $\mathcal{P}$ on which $\mathbf{d}$ is twice continuously differentiable, and define

$$\bar{\mathcal{U}} := \{\bar{\mathbf{p}}(\mathbf{q}) : \mathbf{q} \in \mathcal{U}\} \subset \operatorname{int}(\Delta_{n+1}). \tag{99}$$

The field $\bar{\mathbf{z}}$ is twice continuously differentiable on $\bar{\mathcal{U}}$. Since $\bar{\mathcal{P}}$ is compact, choose a smooth $\chi : \Delta_{n+1} \mapsto [0, 1]$ equal to one on a neighborhood of $\bar{\mathcal{P}}$ and supported in $\bar{\mathcal{U}}$. Then

$$\widetilde{\mathbf{z}}(\bar{\mathbf{p}}) := \begin{cases} \chi(\bar{\mathbf{p}})\bar{\mathbf{z}}(\bar{\mathbf{p}}), & \bar{\mathbf{p}} \in \bar{\mathcal{U}}, \\ \mathbf{0}, & \bar{\mathbf{p}} \notin \bar{\mathcal{U}}, \end{cases} \tag{100}$$

is twice continuously differentiable on Δ_{n+1}, agrees with $\bar{\mathbf{z}}$ near $\bar{\mathcal{P}}$, and satisfies $\langle \bar{\mathbf{p}}, \widetilde{\mathbf{z}}(\bar{\mathbf{p}})\rangle = \chi(\bar{\mathbf{p}})\langle \bar{\mathbf{p}}, \bar{\mathbf{z}}(\bar{\mathbf{p}})\rangle = 0$. ■

F.2 Verification of Theorem 4.3

We need the following result due to Rolf Mantel.

Theorem F.2 (Mantel (1976))**.** *Suppose that $\mathbf{z} : \Delta_n \mapsto \mathcal{X}$ is twice differentiable and satisfies the Walras's law: $\langle \mathbf{p}, \mathbf{z}(\mathbf{p})\rangle = 0$, and let $\mathcal{K} \subset \operatorname{int}(\Delta_n)$ be a compact set. Then there exists an exchange economy with n consumers, in which each consumer's preferences are convex, homothetic, and monotonic, such that the sum of consumer's excess demand function coincides with $\mathbf{z}$ on $\mathcal{K}$.*

Moreover, the endowment directions may be fixed in advance up to a common scale (Mantel, 1976, pp. 198–199): for any linearly independent $\{\boldsymbol{\omega}_t\}_{t\in[n]} \subset \mathbb{R}^n_+$ satisfying $\inf_{\mathbf{p}\in\mathcal{K}}\langle \mathbf{p}, \boldsymbol{\omega}_t\rangle > 0$, the consumers' endowments may be chosen as $\mathbf{b}_t = k\boldsymbol{\omega}_t$ for one sufficiently large finite common $k > 0$.

The above theorem is extension of the Sonnenschein-Mantel-Debreu theorem (Shafer & Sonnenschein, 1982) which says for any continuous function $\mathbf{z}$ satisfying the Walras's law and homogeneity of degree zero, there is an exchange economy whose aggregate excess demand coincides with $\mathbf{z}$ on a compact set. Mantel's result strengthens the result by showing that homothetic agents are sufficient.

Assumption F.2 (Bounded curvature)**.** *For the target excess demand $\mathbf{z}$ and compact price domain $\mathcal{K}$ in Theorem F.2, there exists $\kappa_{\mathbf{z}} < \infty$ such that*

$$\max_{i\in[n]} \sup_{\mathbf{p}\in\mathcal{K}} \left\|\nabla^2 z_i(\mathbf{p})\right\|_{\mathrm{op}} \le \kappa_{\mathbf{z}}. \tag{101}$$

When $\mathbf{z}$ is twice continuously differentiable and $\mathcal{K}$ is compact, Assumption F.2 holds automatically for some $\kappa_{\mathbf{z}}$; stating it separately makes the dependence of the required endowment scale explicit.

Lemma F.1 (Scaled-endowment Mantel representation)**.** *Under the assumptions of Theorem F.2 and Assumption F.2, the representing economy may be chosen so that consumer $t \in [n]$ has endowment*

$$\mathbf{b}_t := (1+\varsigma)\mathbf{e}_t \tag{102}$$

for some finite common scale $1+\varsigma$, where $\varsigma \ge 0$ and $\mathbf{e}_t$ is the t-th coordinate vector of $\mathbb{R}^n$. A sufficient scale depends on $\kappa_{\mathbf{z}}$ and the compact interior domain $\mathcal{K}$. If $\mathbf{x}_t$ denotes consumer t's demand, then

$$\sum_{t\in[n]} (\mathbf{x}_t(\mathbf{p}) - \mathbf{b}_t) = \mathbf{z}(\mathbf{p}), \qquad \mathbf{p}\in\mathcal{K}. \tag{103}$$

Proof. Apply the prescribed-endowment conclusion of Theorem F.2. Choose $\boldsymbol{\omega}_t := \mathbf{e}_t$ for every $t \in [n]$. These vectors are linearly independent, and compactness of $\mathcal{K} \subset \mathrm{int}(\Delta_n)$ implies

$$\inf_{\mathbf{p}\in\mathcal{K}} \langle \mathbf{p}, \mathbf{e}_t\rangle = \inf_{\mathbf{p}\in\mathcal{K}} p_t > 0, \qquad t\in[n], \tag{104}$$

so the prescribed directions satisfy the positive-wealth condition in Theorem F.2. Mantel's construction therefore yields a finite common $k \ge 1$, chosen sufficiently large relative to the derivative bounds of $\mathbf{z}$, including $\kappa_{\mathbf{z}}$, and the domain $\mathcal{K}$, convex, homothetic, and monotonic preferences, and demands $\{\mathbf{x}_t\}_{t\in[n]}$ such that consumer t has endowment $k\mathbf{e}_t$ and

$$\sum_{t\in[n]} (\mathbf{x}_t(\mathbf{p}) - k\mathbf{e}_t) = \mathbf{z}(\mathbf{p}), \qquad \mathbf{p}\in\mathcal{K}. \tag{105}$$

Setting $\varsigma := k - 1 \ge 0$ gives (102) and (103). ■

Together, the lifting construction and Mantel's theorem motivate us to focus on homogeneous real markets, and use homothetic androids with linear wealth.

Theorem F.3. *Let the real market have degree-zero homogeneous aggregate demand satisfying Assumption F.1 and linear wealth functions. Suppose $\mathbf{d}$ is twice continuously differentiable on a neighborhood of a compact price domain $\mathcal{P} \subset \mathbb{R}^n_{++}$ containing the sampled prices, and the associated normalized excess demand satisfies Assumption F.2. For the surrogate market, let $\mathscr{H}$ be the union of all homothetic classes and let the wealth functions be linear. Then, there exist n androids from $\mathscr{H}$ and extra supply $\mathbf{a} = \varsigma\mathbf{1}$, for some $\varsigma \ge 0$, with zero misspecification error in* ($\mathrm{WP}_{+,T}$).

Proof. By homogeneity, it suffices to work on the compact domain of normalized prices,

$$\mathcal{P}_0 := \left\{ \frac{\mathbf{p}}{\langle \mathbf{p}, \mathbf{1}\rangle} : \mathbf{p}\in\mathcal{P} \right\} \subset \mathrm{int}(\Delta_n). \tag{106}$$

By Assumption F.1 and budget exhaustion, $W(\mathbf{p}) = 1$, $\mathbf{g}(\mathbf{p}) = \mathrm{diag}(\mathbf{p})\mathbf{d}(\mathbf{p})$, and $\mathbf{z}(\mathbf{p}) = \mathbf{d}(\mathbf{p}) - \mathbf{1}$ satisfies Walras's law on this domain. Choose an open $\mathcal{V}$ with $\mathcal{P}_0 \subset \mathcal{V}$ and $\overline{\mathcal{V}} \subset \mathrm{int}(\Delta_n)$ on which $\mathbf{z}$ is twice continuously differentiable, and a smooth $\chi : \Delta_n \mapsto [0,1]$ equal to one near $\mathcal{P}_0$ and supported in $\mathcal{V}$. Define $\widetilde{\mathbf{z}} = \chi\mathbf{z}$ on $\mathcal{V}$ and $\widetilde{\mathbf{z}} = \mathbf{0}$ outside $\mathcal{V}$. Then $\widetilde{\mathbf{z}}$ is twice continuously differentiable on Δ_n, agrees with $\mathbf{z}$ near $\mathcal{P}_0$, and satisfies Walras's law. Apply Lemma F.1 to $\widetilde{\mathbf{z}}$ on

$\mathcal{P}_0$. Thus, for some $\varsigma \geq 0$, there are n homothetic androids with $\boldsymbol{\gamma}_t \in \mathscr{H}$ and

$$\mathbf{b}_t = (1+\varsigma)\mathbf{e}_t, \qquad \sum_{t\in[n]} \mathbf{b}_t = (1+\varsigma)\mathbf{1},$$
$$\sum_{t\in[n]} \mathbf{x}_t(\mathbf{p}) - (1+\varsigma)\mathbf{1} = \mathbf{d}(\mathbf{p}) - \mathbf{1}, \qquad \mathbf{p} \in \mathcal{P}_0. \tag{107}$$

Choose $\mathbf{a} = \varsigma\mathbf{1}$ in Section 3.3, so $\sum_t \mathbf{b}_t = \mathbf{1} + \mathbf{a}$. Using $\langle \mathbf{p}, \mathbf{b}_t\rangle \boldsymbol{\gamma}_t(\mathbf{p}) = \operatorname{diag}(\mathbf{p})\mathbf{x}_t(\mathbf{p})$ in (12),

$$\begin{aligned}\mathbf{h}(\mathbf{p}) &= \operatorname{diag}(\mathbf{p}) \sum_{t\in[n]} \mathbf{x}_t(\mathbf{p}) - \varsigma\mathbf{p} \\ &\underset{(107)}{=} \operatorname{diag}(\mathbf{p})\big(\mathbf{d}(\mathbf{p}) + \varsigma\mathbf{1}\big) - \varsigma\mathbf{p} \\ &= \operatorname{diag}(\mathbf{p})\mathbf{d}(\mathbf{p}) \\ &= \mathbf{g}(\mathbf{p}), \qquad \mathbf{p} \in \mathcal{P}_0. \end{aligned} \tag{108}$$

Every sampled loss in ($\mathrm{WP}_{+,T}$) therefore vanishes, proving zero misspecification error. ■

The same argument applied to the lifted excess demand in (96) uses its supply $[\mathbf{1}; M]$ in the correction and recovers the original shares by

$$\bar{\mathbf{h}}(\bar{\mathbf{p}}) := \operatorname{diag}(\bar{\mathbf{p}})[\mathbf{1}; M] + \sum_{t\in[n+1]} \left[\langle \bar{\mathbf{p}}, \mathbf{b}_t\rangle \boldsymbol{\gamma}_t(\bar{\mathbf{p}}) - \operatorname{diag}(\bar{\mathbf{p}})\mathbf{b}_t\right],$$
$$\mathbf{h}(\mathbf{q}) := \frac{\bar{\mathbf{h}}_{1:n}(\bar{\mathbf{p}}(\mathbf{q}))}{\langle \mathbf{1}, \bar{\mathbf{h}}_{1:n}(\bar{\mathbf{p}}(\mathbf{q}))\rangle}. \tag{109}$$

The same excess-demand identity implies $\bar{\mathbf{h}}(\bar{\mathbf{p}}) = \operatorname{diag}(\bar{\mathbf{p}})\bar{\mathbf{d}}(\bar{\mathbf{p}})$, so we have $\frac{\operatorname{diag}(\mathbf{q})\mathbf{d}(\mathbf{q})}{W(\mathbf{q})} = \mathbf{g}(\mathbf{q})$ whenever $W(\mathbf{q}) > 0$.

Remark 4. *On the fixed compact interior price domain $\mathcal{P}$, there is $R < \infty$ such that $\|\mathbf{p}\| \leq R$ for every $\mathbf{p} \in \mathcal{P}$. Together with the C^2 assumption in Theorem F.3, this ensures that Mantel's construction admits a finite $\varsigma \geq 0$, chosen independently of the sample. In particular, the common scale $1+\varsigma$ is chosen above a finite threshold determined by the curvature bound $\kappa_{\mathbf{z}}$ and the compact interior price domain. Its required size depends on the excess demand and its derivatives as well as the price domain; R alone does not provide a uniform bound. For any fixed budget $D_{\mathbf{a}} \geq \varsigma$, the exact representation with $\mathbf{a} = \varsigma\mathbf{1}$ belongs to the class in Theorem 4.1, provided $\mathscr{H}$ contains the representing androids.*

G Proof of results on misspecification error II

We prove Theorem G.1, the formal version of Theorem 4.4, by an example with $n = 2$ goods. Because of that, we can focus on the expenditure share of the first good (i.e., $g_1(\mathbf{p})$). We consider androids that are *simple* if its share $\gamma_1(\mathbf{p})$ has bounded total variation, and write $\mathrm{TV}(\boldsymbol{\gamma}_1) \leq G < \infty$. The proof proceeds in four steps:

(1) We first parameterize the two-good economy by the log-ratio of prices τ and recall the properties of total variation (Section G.1, Lemma G.1).

(2) We then show that, if every android's wealth is a polynomial of degree q in the Bernstein basis, the total variation of the surrogate share is at most $q + (q+1)G$, independently of the number of androids (Section G.2, Lemma G.3). The Fisher ($q = 0$) and Arrow-Debreu ($q = 1$) markets are special cases.

(3) Next, we construct a homothetic agent whose expenditure share oscillates in τ and thus has unbounded total variation (Section G.3, Example 2).

(4) Finally, a share of bounded total variation cannot uniformly track a share that oscillates infinitely often (Lemma G.4). This yields the lower bound $\frac{\tanh(1)}{2}$, which is attained by Cobb-Douglas androids (Section G.4, Theorem G.1).

G.1 TWO-GOOD ECONOMY AND TOTAL VARIATION

For a homothetic agent with with two goods, we can write the quantities through the log-ratio of prices. Consider the quantity ratio r_x and the expenditure-share ratio r_γ as function of the log-ratio of prices,

$$\tau(\mathbf{p}) := \log \frac{p_1}{p_2}, \quad r_x(\tau) := \frac{x_1}{x_2}, \quad r_\gamma(\tau) := \frac{\gamma_1(\tau)}{\gamma_2(\tau)} = \exp(\tau) r_x(\tau). \tag{110}$$

Write $\ell(z) := \frac{\exp(z)}{1+\exp(z)}$ for the logistic map, so that $\gamma_1 = \ell(\log r_\gamma)$. We can define the elasticity of substitution is defined locally by

$$\frac{\mathrm{d} \log r_\gamma(\tau)}{\mathrm{d}\tau} = \frac{\gamma_1'(\tau)}{\gamma_1(1-\gamma_1)} = -\sigma(\tau), \qquad 1 + \sigma(\tau) \geq 0,$$

where the logistic factor uses $\gamma_1 = \ell(\log r_\gamma)$.

For a function $f : \mathcal{I} \mapsto \mathbb{R}$ on an interval $\mathcal{I} \subseteq \mathbb{R}$, let $\mathrm{TV}_\mathcal{I}(f)$ be the total variation of f on $\mathcal{I}$, defined as the supremum of the sum of the absolute differences of the function over all possible partitions of the interval. Namely, for a finite subset $S_m = \{\tau_0, \cdots, \tau_m\} \subseteq \mathcal{I}$ with $\tau_0 < \cdots < \tau_m \in \mathcal{I}$,

$$\mathrm{TV}_\mathcal{I}(f) = \sup_{S_m} \sum_{j=1}^{m} |f(\tau_j) - f(\tau_{j-1})|.$$

We write $\mathrm{TV}(f) := \mathrm{TV}_\mathbb{R}(f)$ for short. The following properties are standard; see Folland (1999, §3.5).

Lemma G.1 (Total variation)**.** *The following properties hold for the total variation of a function $f : \mathcal{I} \mapsto \mathbb{R}$ on an interval $\mathcal{I} \subseteq \mathbb{R}$:*

(1) *If f is absolutely continuous, then $\mathrm{TV}_\mathcal{I}(f) = \int_\mathcal{I} |f'(\tau)| \mathrm{d}\tau$.*

(2) *If f is monotone, then $\mathrm{TV}_\mathcal{I}(f) = \sup_\mathcal{I} f - \inf_\mathcal{I} f$; if further f is valued in $[0,1]$, then $\mathrm{TV}_\mathcal{I}(f) \leq 1$.*

(3) *$\mathrm{TV}(af) = |a| \mathrm{TV}(f)$, $\mathrm{TV}\left(\sum_k \lambda_k f_k\right) \leq \sum_k \lambda_k \mathrm{TV}(f_k)$ for $\lambda_k \geq 0$.*

(4) *If $f_m \to f$ pointwise, then $\mathrm{TV}(f) \leq \liminf_m \mathrm{TV}(f_m)$.*

Remark 5. *A CES agent has bounded TV and $G = 1$: its log-odds is affine in τ, so its share is monotone, hence $\mathrm{TV} \leq 1$ by Lemma G.1(2). A linear agent also has finite variation: in two goods, linear utility $u(\mathbf{x}) = a_1 x_1 + a_2 x_2$ with $a_1, a_2 > 0$ yields*

$$\gamma_1(\tau) = \begin{cases} 1, & \tau \leq \log \frac{a_1}{a_2}, \\ 0, & \tau > \log \frac{a_1}{a_2}, \end{cases}$$

with the tie value arbitrary, so $\mathrm{TV}(\gamma_1) = 1$.

G.2 THE ROLE OF WEALTH FUNCTIONS

Now consider the wealth function $w_t(\mathbf{p})$ of the android t. We focus on polynomial wealth functions: each android's wealth w_t is a nonnegative combination of the degree-q Bernstein basis in the price of the first good $p_1 = \ell(\tau)$ (a subclass of the nonnegative degree-q polynomials). That is, $w_t \in \mathscr{B}^q$, defined as follows.

$$\hat{w}_k(\tau) := B_k^q(p_1(\tau)), \qquad B_k^q(s) := \binom{q}{k} s^k (1-s)^{q-k}, \quad k = 0, \ldots, q,$$

$$w_t(\mathbf{p}) := \sum_{k=0}^{q} c_{t,k} \hat{w}_k(\tau), \qquad c_{t,k} \geq 0. \tag{$\mathscr{B}^q$}$$

We keep $\sum_t^T c_{t,k} = 1$ for every k, so the expenditure share regroups from a sum over the T androids into a sum over the $q+1$ weights,

$$h_1 = \sum_{t=1}^{T} w_t(\mathbf{p}) \gamma_{t1} = \sum_{k=0}^{q} \hat{w}_k(\tau) \hat{g}_{k1}, \qquad \hat{g}_{k1} := \sum_{t=1}^{T} c_{t,k} \gamma_{t1}. \tag{111}$$

For each k, $\hat{g}_{k1}$ is a convex combination of the base shares γ_{t1} with the constant weights $c_{t,k}$; by Lemma G.1(3), we immediately have

$$\mathrm{TV}(\hat{g}_{k1}) \le G.$$

We need the following facts about the Bernstein basis.

Lemma G.2 (Bernstein basis)**.** *The following holds for the degree-q Bernstein basis $B_k^q(s)$ on $[0,1]$:*

(1) $\sum_{k=0}^{q} B_k^q(s) = 1$.

(2) $(B_k^q)'(s) = q\left(B_{k-1}^{q-1}(s) - B_k^{q-1}(s)\right)$*, with out-of-range terms zero.*

(3) $\int_0^1 \sum_{k=0}^{q} |(B_k^q)'(s)| \mathrm{d}s \le 2q$.

Proof. Statements (1) and (2) are standard. For (3), we note

$$\sum_{k=0}^{q} |(B_k^q)'(s)| \le q \sum_{k=0}^{q} \left(B_{k-1}^{q-1}(s) + B_k^{q-1}(s)\right) = 2q,$$

and integrating over $[0,1]$ gives the bound. ∎

Lemma G.3 (q-th order additive wealth shares)**.** *Let $n = 2$. Any h_1 following (111) satisfies*

$$\mathrm{TV}(h_1) \le q + (q+1)G,$$

independently of the number of androids.

Proof. We argue directly from the partition definition of total variation, so the base shares $\hat{g}_{k1}$ need only be of bounded variation (the weights $\hat{w}_k$ are smooth). Fix a partition $S_m = \{\tau_0 < \cdots < \tau_m\}$ and write $\Delta_j f := f(\tau_j) - f(\tau_{j-1})$. Since $\sum_k \hat{w}_k \equiv 1$ by Lemma G.2(1), we have $\sum_k \Delta_j \hat{w}_k = 0$. Summation by parts on (111) gives

$$\begin{aligned}\Delta_j h_1 &= \sum_{k=0}^{q} \left[\hat{w}_k(\tau_j)\,\Delta_j \hat{g}_{k1} + \hat{g}_{k1}(\tau_{j-1})\,\Delta_j \hat{w}_k\right] \\ &= \sum_{k=0}^{q} \hat{w}_k(\tau_j)\,\Delta_j \hat{g}_{k1} + \sum_{k=0}^{q} \left(\hat{g}_{k1}(\tau_{j-1}) - \tfrac{1}{2}\right)\Delta_j \hat{w}_k,\end{aligned}$$

the second line subtracting $\frac{1}{2}\sum_k \Delta_j \hat{w}_k = 0$. Since $0 \le \hat{w}_k \le 1$ and $|\hat{g}_{k1} - \frac{1}{2}| \le \frac{1}{2}$,

$$|\Delta_j h_1| \le \sum_{k=0}^{q} |\Delta_j \hat{g}_{k1}| + \tfrac{1}{2}\sum_{k=0}^{q} |\Delta_j \hat{w}_k|.$$

Summing over j and taking the supremum over partitions,

$$\mathrm{TV}(h_1) \le \sum_{k=0}^{q} \mathrm{TV}(\hat{g}_{k1}) + \tfrac{1}{2}\sum_{k=0}^{q} \mathrm{TV}(\hat{w}_k).$$

Each $\hat{g}_{k1}$ satisfies $\mathrm{TV}(\hat{g}_{k1}) \le G$, so the first sum is at most $(q+1)G$. The weights $\hat{w}_k = B_k^q(p_1(\tau))$ are absolutely continuous, and since $p_1 = \ell(\tau)$ is monotone onto $(0,1)$, the change of variables $s = p_1$ with Lemma G.2(3) gives

$$\sum_{k=0}^{q} \mathrm{TV}(\hat{w}_k) = \int_{\mathbb{R}} \sum_{k=0}^{q} |\hat{w}_k'(\tau)| \mathrm{d}\tau = \int_0^1 \sum_{k=0}^{q} |(B_k^q)'(s)| \mathrm{d}s \le 2q.$$

Therefore $\mathrm{TV}(h_1) \le (q+1)G + q$. ∎

The Fisher and Arrow-Debreu markets are only special cases.

- $q = 0$ *(Fisher).* The only basis function is $B_0^0 \equiv 1$, so $\hat{w}_0 \equiv 1$ and the wealth

  $$w_t = c_{t,0}$$

  is a price-independent constant (fixed budget). Then $h_1 = \sum_t c_{t,0}\gamma_{t1}$, giving $\mathrm{TV}(h_1) \le G$.

- $q = 1$ *(Arrow-Debreu).* The basis is $B_0^1(s) = 1 - s$ and $B_1^1(s) = s$, so with $s = p_1$ and $p_1 + p_2 = 1$ on Δ_2 the weights are $\hat{w}_0 = p_2$, $\hat{w}_1 = p_1$, and the wealth
$$w_t = c_{t,0}p_2 + c_{t,1}p_1 = \langle \mathbf{p}, \mathbf{b}_t \rangle, \qquad \mathbf{b}_t = (c_{t,1}, c_{t,0}),$$
is the value of the endowment $\mathbf{b}_t$. Then $h_1 = \sum_t \langle \mathbf{p}, \mathbf{b}_t \rangle \gamma_{t1}$, giving the Arrow-Debreu bound $\mathrm{TV}(h_1) \le 2G + 1$.

G.3 A HOMOTHETIC AGENT WITH UNBOUNDED VARIATION

In the following one-agent economy, the market share $\mathbf{g}$ is the agent's own share $\boldsymbol{\gamma}$. We present its share versus the price of the first good p_1 in Figure 1.

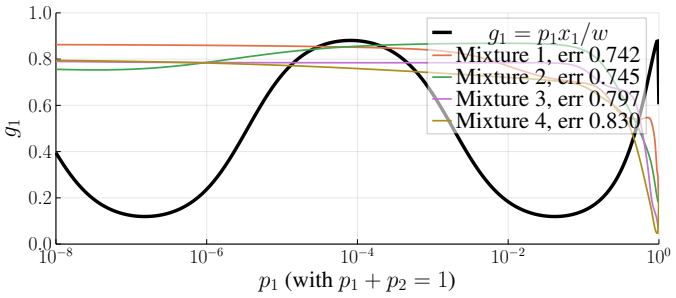

Figure 1: Expenditure share of Example 2. Four different random mixtures of CES agents are presented for comparison.

Example 2 (A homothetic "oscillating" agent)**.** *Let* $r_x(\tau) := \exp\left(-\tau + 2\sin(\frac{\tau}{2})\right)$. *Then it can be rationalized by the utility*

$$u(x_1, x_2) := x_2 \phi\left(\frac{x_1}{x_2}\right), \quad \textit{where} \quad \log \phi(\rho) := \int_1^\rho \frac{\mathrm{d}s}{s + \exp(-r_x^{-1}(s))} < \infty. \tag{112}$$

Furthermore, it holds that

(1) u *is homogeneous of degree one, increasing, and concave.*

(2) *The expenditure share has log-odds*
$$\log \frac{g_1(\tau)}{g_2(\tau)} = 2\sin(\tfrac{\tau}{2}), \tag{113}$$
hence $\mathrm{TV}(g_1) = \infty$.

Proof. We first verify that the integral is well-defined. The exponent $\log r_x(\tau) = -\tau + 2\sin(\frac{\tau}{2})$ has derivative $\cos(\frac{\tau}{2}) - 1 \le 0$, so r_x is a strictly decreasing continuous bijection from $\mathbb{R}$ onto $(0, \infty)$, and r_x^{-1} is defined and continuous on $(0, \infty)$. The integrand is continuous on $(0, \infty)$ and satisfies $0 < [s + \exp(-r_x^{-1}(s))]^{-1} \le \frac{1}{s}$. Hence for any $\rho > 0$,
$$|\log \phi(\rho)| = \left| \int_1^\rho \frac{\mathrm{d}s}{s + \exp(-r_x^{-1}(s))} \right| \le \left| \int_1^\rho \frac{\mathrm{d}s}{s} \right| = |\log \rho| < \infty,$$
so the integral in (112) is finite and $\log \phi(\rho)$ is well-defined.

For part (1), homogeneity is immediate from (112). Let
$$\rho := \frac{x_1}{x_2}, \quad D(\rho) := \rho + \exp(-r_x^{-1}(\rho)) > \rho,$$
so that (112) reads $\frac{\phi'(\rho)}{\phi(\rho)} = \frac{1}{D(\rho)}$. With $\mathrm{d}_{x_1}\rho = 1/x_2$ and $\mathrm{d}_{x_2}\rho = -\rho/x_2$, differentiating $u = x_2\phi(\rho)$ gives
$$u_1' = x_2\phi'(\rho)\mathrm{d}_{x_1}\rho = \phi'(\rho),$$
$$u_2' = \phi(\rho) + x_2\phi'(\rho)\mathrm{d}_{x_2}\rho = \phi(\rho) - \rho\phi'(\rho) = \phi(\rho)\frac{D(\rho) - \rho}{D(\rho)},$$

both positive (since $\phi > 0$ and $D(\rho) > \rho$), so u is increasing, with marginal rate of substitution

$$\frac{u_1'}{u_2'} = \frac{\phi'(\rho)}{\phi(\rho) - \rho\phi'(\rho)} = \exp(r_x^{-1}(\rho)). \tag{114}$$

For concavity we show ϕ is concave; then $u(x_1, x_2) = x_2\phi(x_1/x_2)$, its perspective, is concave on $\{x_2 > 0\}$. Since r_x^{-1} is decreasing, $\exp(-r_x^{-1}(\rho))$ is increasing, so $D(\rho) - \rho = \exp(-r_x^{-1}(\rho))$ is increasing and $\mathrm{d}D \geq \mathrm{d}\rho$ as Stieltjes measures. From $\phi' = \phi/D$ we have $\log \phi' = \log \phi - \log D$, so for $0 < a < b$, using $(\log \phi)' = 1/D$ and $\mathrm{d} \log D = \mathrm{d}D/D$,

$$\log \frac{\phi'(b)}{\phi'(a)} = \int_a^b \frac{\mathrm{d}\rho}{D(\rho)} - \int_{(a,b]} \frac{\mathrm{d}D(\rho)}{D(\rho)} \leq 0,$$

since $\mathrm{d}D \geq \mathrm{d}\rho$ and $D > 0$. Hence ϕ' is nonincreasing, ϕ is concave, and so is u. Part (1) is proved.

Since u is increasing, the budget binds, and the KKT condition for utility maximization reads

$$-\nabla u + \lambda \mathbf{p} = 0 \implies \frac{u_1'}{u_2'} = \frac{p_1}{p_2} = \exp(\tau).$$

By (114) this is $\exp(r_x^{-1}(\rho)) = \exp(\tau)$, so $\rho = r_x(\tau)$ and the optimizer is

$$x_1(\mathbf{p}, w) = \frac{w r_x(\tau)}{p_1 r_x(\tau) + p_2}, \quad x_2(\mathbf{p}, w) = \frac{w}{p_1 r_x(\tau) + p_2}.$$

Then

$$\frac{g_1(\tau)}{g_2(\tau)} = \frac{p_1 x_1}{p_2 x_2} = \exp(\tau) r_x(\tau) = \exp\left(2\sin(\tfrac{\tau}{2})\right),$$

which gives the log-odds (113). Inverting the logistic map,

$$g_1(\tau) = \ell\left(2\sin(\tfrac{\tau}{2})\right).$$

This share oscillates between $\ell(-2)$ and $\ell(2)$ every 4π units of log-price. On the partition S_{2N} with $\tau_j = \pi + 2\pi j$, $j = 0, \ldots, 2N$, the identity $2\sin(\frac{\tau_j}{2}) = 2(-1)^j$ gives $g_1(\tau_j) = \ell(2(-1)^j)$, so consecutive values alternate between $\ell(2)$ and $\ell(-2)$ and

$$\sum_{j=1}^{2N} |g_1(\tau_j) - g_1(\tau_{j-1})| = 2N\left(\ell(2) - \ell(-2)\right),$$

which diverges as $N \to \infty$. Hence $\mathrm{TV}(g_1) = \infty$. ■

Economically, the elasticity in Example 2 is $\sigma(\tau) = -\cos(\frac{\tau}{2})$, oscillating over $[-1, 1]$. As p_1 rises the share g_1 alternately rises and falls, so the first good is neither a global complement nor a global substitute.

G.4 A SHARP LOWER BOUND

Lemma G.4. *Let $h : \mathbb{R} \mapsto [0, 1]$ with $\mathrm{TV}(h) \leq C$. Then for every partition $S_m = \{\tau_0 < \cdots < \tau_m\}$ and fixed $f : \mathbb{R} \mapsto [0, 1]$,*

$$\sup_{\tau \in \mathbb{R}} |f(\tau) - h(\tau)| \geq \frac{1}{2m} \max\left\{ \sum_{j=1}^{m} |f(\tau_j) - f(\tau_{j-1})| - C,\ 0 \right\}. \tag{115}$$

Proof. Bounding each pointwise gap $|f(\tau_i) - h(\tau_i)|$ by $\sup_{\tau \in \mathbb{R}} |f(\tau) - h(\tau)|$, for each adjacent pair,

$$\begin{aligned} |f(\tau_j) - f(\tau_{j-1})| &\leq |h(\tau_j) - h(\tau_{j-1})| + |f(\tau_j) - h(\tau_j)| + |f(\tau_{j-1}) - h(\tau_{j-1})| \\ &\leq |h(\tau_j) - h(\tau_{j-1})| + 2\sup_{\tau \in \mathbb{R}} |f(\tau) - h(\tau)|. \end{aligned}$$

Summing over S_m gives

$$\begin{aligned}\sum_{j=1}^{m} |f(\tau_j) - f(\tau_{j-1})| &\le \sum_{j=1}^{m} |h(\tau_j) - h(\tau_{j-1})| + 2m \sup_{\tau\in\mathbb{R}} |f(\tau) - h(\tau)| \\ &\le \mathrm{TV}(h) + 2m \sup_{\tau\in\mathbb{R}} |f(\tau) - h(\tau)| \\ &\le C + 2m \sup_{\tau\in\mathbb{R}} |f(\tau) - h(\tau)|.\end{aligned}$$

Rearranging gives (115); if the bracketed term is negative, the bound holds trivially since $\sup_{\tau\in\mathbb{R}} |f(\tau) - h(\tau)| \ge 0$. ■

Theorem G.1 (Sharp lower bound for fixed finite order)**.** *Fix $q < \infty$ and $1 \le T \le \infty$, and let g_1 be the target of Example 2. The best uniform approximation over admissible pairs of wealth functions and base shares is*

$$\inf_{\substack{w_t\in\mathscr{B}^q,\ \sum_t w_t\equiv 1 \\ \gamma_{t1}:\mathbb{R}\mapsto[0,1],\ \mathrm{TV}(\gamma_{t1})\le G}} \sup_{\tau\in\mathbb{R}} \left| g_1(\tau) - \sum_{t=1}^{T} w_t(\mathbf{p})\gamma_{t1}(\tau) \right| = \frac{\ell(2) - \ell(-2)}{2} = \frac{\tanh(1)}{2} \approx 0.3808,$$

The bound is attained if every android t is uses a Cobb-Douglas share.

Proof. Take the partition S_m with $\tau_j := \pi + 2\pi j,\ j = 0, \dots, m$. Then $\sin(\tau_j/2) = (-1)^j$, so for any j,

$$|g_1(\tau_j) - g_1(\tau_{j-1})| = \ell(2) - \ell(-2).$$

Since $w_t \in \mathscr{B}^q$ for every t, Lemma G.3 gives $\mathrm{TV}(h_1) \le q + (q+1)G$. Applying Lemma G.4 to S_m with $C = q + (q+1)G$ gives

$$\sup_{\tau\in\mathbb{R}} |g_1(\tau) - h_1(\tau)| \ge \frac{m(\ell(2) - \ell(-2)) - (q + (q+1)G)}{2m}.$$

Taking $m \to \infty$ yields

$$\sup_{\tau\in\mathbb{R}} |g_1(\tau) - h_1(\tau)| \ge \frac{\ell(2) - \ell(-2)}{2}.$$

Finally,

$$\ell(2) - \ell(-2) = \frac{\exp(2) - 1}{\exp(2) + 1} = \tanh(1),$$

and $\ell(-2) = 1 - \ell(2)$, so the constant share $h_1 \equiv \frac{1}{2}$ has uniform error $(\ell(2) - \ell(-2))/2$. This constant is attainable by taking every android Cobb-Douglas such that $\gamma_{t1} \equiv \frac{1}{2}$, and thus,

$$\gamma_{t1}(\tau) \equiv \tfrac{1}{2}, \quad h_1(\tau) \quad = \sum_t w_t(\mathbf{p})\gamma_{t1}(\tau) = \tfrac{1}{2}\sum_t w_t(\mathbf{p}) = \tfrac{1}{2}.$$

Since $2\sin(\frac{\tau}{2})$ ranges over $[-2, 2]$ and ℓ is increasing with $\ell(-2) = 1 - \ell(2)$, the share $g_1(\tau) = \ell(2\sin(\frac{\tau}{2}))$ is farthest from $\frac{1}{2}$ at the extremes, so

$$\sup_{\tau\in\mathbb{R}} |g_1(\tau) - h_1(\tau)| = \sup_{\tau\in\mathbb{R}} \left|\ell\left(2\sin(\tfrac{\tau}{2})\right) - \tfrac{1}{2}\right| = \ell(2) - \tfrac{1}{2} = \frac{\ell(2) - \ell(-2)}{2} = \frac{\tanh(1)}{2}.$$

This matches the lower bound established above. ■

Remark 6. *The aggregate demand of homothetic androids with fixed budgets is the gradient of a convex potential function (Eisenberg & Gale, 1959), so the surrogate market cannot match any case that fails this integrability restriction (such as those from non-homothetic agents), no matter how many androids are introduced. This difficulty can be partly alleviated by using non-homothetic androids while not needed in theory.*

H METHODS OF SOLVING THE SEPARATION ORACLE

For the convenience of our presentation, we drop the subscript T in this section. Let us recall the separation oracle (SEP) below needed for finding new android at each iteration:

$$\text{SEP}: \pi^* = \max_{\boldsymbol{\gamma} \in \mathscr{H}} \sum_{k=1}^{K} \langle \mathbf{V}\mathbf{e}_k, \boldsymbol{\gamma}(\mathbf{p}_k) \rangle. \tag{116}$$

It can be written using the parameterization via $(\mathbf{y}, \sigma)$; see Section B. The following fact about SEP will be repeatedly used in this section.

Fact 7. *For any $a \in \mathbb{R}$, SEP is invariant under affine addition, viz.,*

$$\pi^* = \max_{\boldsymbol{\gamma} \in \mathscr{H}} \sum_{k=1}^{K} \langle \mathbf{V}\mathbf{e}_k, \boldsymbol{\gamma}(\mathbf{p}_k) \rangle = \max_{\boldsymbol{\gamma} \in \mathscr{H}} \sum_{k=1}^{K} \langle \mathbf{V}\mathbf{e}_k + a\mathbf{1}, \boldsymbol{\gamma}(\mathbf{p}_k) \rangle - a \cdot K.$$

Hence, we can always assume that $\mathbf{V}\mathbf{e}_k \geq \mathbf{0}$ for all $k \in [K]$. In the following, we show that SEP splits into two categories:

(1) The linear case $r = 1, \sigma = \infty$, cf. Section H.1. With log-price data, this case is polynomial-time solvable for every fixed K and strongly NP-hard when K is part of the input, by a reduction from Maximum Acyclic Subgraph (MAS).

(2) The remaining cases, including the CES case $\sigma \in (-1, \infty), r \in (-\infty, 1)$ (including the Cobb-Douglas class $\sigma = 0$) and the Leontief case $\sigma = -1, r = -\infty$. We show the reduction to a Linear Fractional Programming (LFP) problem in Section H.2 and a FPTAS for global optimal solution in Section H.3.

For the first category, a mixed-integer solver generates candidate choices, and a linear-programming margin test enforces the fixed tie-breaking rule through exclusion cuts; see Section H.1.

Because the practical difficulty of solving SEP globally in the latter category, we show an approximation to KKT point is easy to find (cf. Section H.4). In theory, this also allows a FPTAS (cf. Remark 7), and off-the-shelf interior point solvers can be used.

H.1 LINEAR CASE

For linear utilities, the fixed smallest-index rule in Fact 5 selects

$$\boldsymbol{\gamma}(\mathbf{p}_k) = \mathbf{e}_{j_k^*}, \qquad j_k^* = \min \arg\max_{j \in [n]} \{y_j - \log p_{k,j}\}, \quad k \in [K].$$

Weak maximizing inequalities alone allow inconsistent choices at ties. For example, if $\mathbf{p}_1 = \mathbf{p}_2 = (\frac{1}{2}, \frac{1}{2})$, $\mathbf{V}\mathbf{e}_1 = \mathbf{e}_1$, and $\mathbf{V}\mathbf{e}_2 = \mathbf{e}_2$, deterministic demand has objective value 1. The weak inequalities admit $\mathbf{y} = \mathbf{0}$, $\boldsymbol{\gamma}_1 = \mathbf{e}_1$, and $\boldsymbol{\gamma}_2 = \mathbf{e}_2$, with objective value 2. To enforce deterministic demand, define the incompatible choice patterns using (26):

$$\mathcal{J}_{\text{bad}} := \{\mathbf{j} \in [n]^K : \eta^*(\mathbf{j}) \leq 0\}. \tag{117}$$

For positive simplex prices and the unrestricted parameterization of $\mathscr{H}^{\text{Lin}}$, an exact formulation of SEP is

$$\begin{aligned}
\pi^* = \max_{\boldsymbol{\gamma}_1, \ldots, \boldsymbol{\gamma}_K, \mathbf{y}} \; & \sum_{k=1}^{K} \langle \mathbf{V}\mathbf{e}_k, \boldsymbol{\gamma}_k \rangle \\
\text{s.t.} \; & \boldsymbol{\gamma}_k \in \{0,1\}^n, \; \langle \mathbf{1}, \boldsymbol{\gamma}_k \rangle = 1, && k \in [K], \\
& y_j - \log p_{k,j} \geq y_\ell - \log p_{k,\ell} - M(1 - \gamma_{k,j}), && j, \ell \in [n], \; k \in [K], \qquad \text{SEP-}\mathscr{H}^{\text{Lin}} \\
& \sum_{k=1}^{K} \gamma_{k,j_k} \leq K - 1, && \mathbf{j} \in \mathcal{J}_{\text{bad}}, \\
& y_n = 0.
\end{aligned}$$

Each exclusion cut removes precisely one incompatible pattern. The big-M constant can be chosen as

$$D := \max_{k,j} |\log p_{k,j}|, \qquad M \geq 2D.$$

Indeed, shift any weak maximizing witness so that $\max_j y_j = 0$, and raise coordinates below $-D$ to $-D$. At each sample, some unmodified coordinate has score at least zero, while every raised coordinate has score at most zero, so the selected weak maximizers remain maximizing. All resulting scores lie in $[-D, D]$; a final common shift enforces $y_n = 0$ without changing score differences. Thus the big-M inequalities retain every realizable choice pattern. Their auxiliary vector $\mathbf{y}$ certifies weak maximality; a parameter that realizes the pattern under the fixed rule is recovered from (26) with positive margin.

The exclusion cuts can be generated as needed. Start with the mixed-integer program without these cuts and solve it to global optimality. For its selected pattern $\mathbf{j}$, solve (26). If $\eta^*(\mathbf{j}) > 0$, return the parameter from this margin program. Otherwise, add the cut $\sum_{k=1}^{K} \gamma_{k,j_k} \le K - 1$ and solve the mixed-integer program again. Every intermediate mixed-integer optimum is an upper bound on the deterministic oracle, since only incompatible patterns are excluded. An accepted pattern attains that bound with a single deterministic demand function, proving exactness. The procedure terminates after examining at most n^K patterns. The connection between SEP-$\mathscr{H}^{\mathrm{Lin}}$ and the Maximum Acyclic Subgraph (MAS) problem uses the following definition.

Definition H.1 (MAS). *Given a directed graph $G = (V, E)$, the MAS is to find the largest acyclic edge subset $S \subseteq E$.*

MAS is equivalent to minimum feedback edge set and is strongly NP-hard (Karp, 1972). For the following complexity statement, the input consists of rational log-price representatives and rational rewards; normalizing each price vector does not change the selected good.

Theorem H.1. *SEP for linear utility* (SEP-$\mathscr{H}^{\mathrm{Lin}}$) *is polynomial-time solvable for every fixed K and* strongly *NP-hard when K is part of the input.*

Proof.

Case 1: fixed K. Enumerate the n^K choice patterns $\mathbf{j} \in [n]^K$ and solve (26) for each pattern. This linear program has $n + 1$ variables and $\mathcal{O}(Kn)$ constraints, and the pattern is realizable under the fixed tie-breaking rule if and only if $\eta^*(\mathbf{j}) > 0$. Among the realizable patterns, return one maximizing $\sum_{k=1}^{K} \langle \mathbf{V}\mathbf{e}_k, \mathbf{e}_{j_k} \rangle$, together with its realizing parameter. Each linear program has polynomial input length, so enumeration takes polynomial time for every fixed K. For $K = 1$, every good is realizable by Fact 6, and scanning the n rewards suffices.

Case 2: K is part of the input. We establish strong NP-hardness by a reduction from the maximum acyclic subgraph (MAS) problem to the deterministic linear oracle, constructing only its price and reward data.

Consider a graph $G = (V, E)$. Index the edges as $E = \{e_1, \ldots, e_{|E|}\}$ and create one sample per edge, so $K = |E|$ and sample k corresponds to e_k; write $e_k = (a_k, b_k)$, where $a_k \in V, b_k \in V$. The goods are one per vertex of V, plus one extra *skip good* d_k used only by sample k. Fix integers $Y > 2K + 1$ and $B > 2Y + 1$. For each sample k, and the edge e_k, set the data $(\mathbf{p}_k, \mathbf{u}_k)$ as

$$\begin{aligned} &\log p_{k,a_k} = 0, \quad \log p_{k,b_k} = -1, \quad \log p_{k,d_k} = 0, \quad \log p_{k,\ell} = B \ \ (\ell \neq a_k, b_k, d_k), \\ &u_{k,a_k} = 1, \quad u_{k,\ell} = 0 \ \ (\ell \neq a_k). \end{aligned} \tag{118}$$

The displayed prices can be normalized to the simplex without changing any score differences; we compute scores using the displayed representatives. A feasible solution of (SEP-$\mathscr{H}^{\mathrm{Lin}}$) with data (118) chooses a good j_k^* for each sample k. If $j_k^* = a_k$, it obtains reward $u_{k,a_k} = 1$.

Now, let us show that $\pi^* \le \mathrm{MAS}(G)$. Take any feasible $(\boldsymbol{\gamma}, \mathbf{y})$ and let $S = \{k : \gamma_{k,a_k} = 1\}$. It is apparent that the objective value equals to $|S|$. For each $k \in S$, let $j' = b_k$ gives

$$y_{a_k} - y_{b_k} \ge 1.$$

Summing the above inequalities over any directed cycle $C \subseteq S$, the left side telescopes to 0 while the right side is $|C| \ge 1$, a contradiction. Hence $S \subseteq E$ must be cycle free, so $|S| \le \mathrm{MAS}(G)$.

Conversely, let $S \subseteq E$ be acyclic. Let ℓ_v be the number of edges on the longest directed path of (V, S) starting at v, and choose

$$y_v = 2\ell_v \in [0, 2K], \quad v \in V, \qquad y_{d_k} = \begin{cases} -1, & k \in S, \\ Y, & k \notin S. \end{cases}$$

For $k \in S$, good a_k uniquely maximizes the score:

$$\begin{aligned}(y_{a_k} - \log p_{k,a_k}) - (y_{b_k} - \log p_{k,b_k}) &= y_{a_k} - y_{b_k} - 1 \geq 1,\\ (y_{a_k} - \log p_{k,a_k}) - (y_{d_k} - \log p_{k,d_k}) &= y_{a_k} + 1 \geq 1,\\ (y_{a_k} - \log p_{k,a_k}) - (y_v - \log p_{k,v}) &\geq B - 2K > 0, \quad v \in V \setminus \{a_k, b_k\},\\ (y_{a_k} - \log p_{k,a_k}) - (y_{d_h} - \log p_{k,d_h}) &\geq B - Y > 0, \quad h \neq k.\end{aligned}$$

For $k \notin S$, its own skip good d_k has score Y. The scores of a_k and b_k are at most $2K$ and $2K+1$, respectively, and every other good has score at most $Y - B$. Hence d_k is also the unique maximizer. These strict choices obey the fixed tie-breaking rule and have positive common margin in (26), so none is excluded by the cuts in (SEP-$\mathscr{H}^{\mathrm{Lin}}$). The score differences have magnitude at most $B + Y + 1 < 2B$, so the big-M inequalities also hold. Subtracting y_n from every parameter enforces the gauge without changing the choices. The resulting objective equals $|S|$, and choosing a maximum acyclic subgraph establishes $\pi^* \geq \mathrm{MAS}(G)$. Together, we have proved that $\pi^* = \mathrm{MAS}(G)$. The reduction is exact in objective value.

The construction has $K = |E|$ and runs in time polynomial in $|V| + |E|$. Its numerical input data are polynomially bounded: $Y, B = \mathcal{O}(|E|)$, the log-price representatives lie in $\{-1, 0, B\}$, and the rewards lie in $\{0, 1\}$. By the strong NP-hardness of MAS, SEP-$\mathscr{H}^{\mathrm{Lin}}$ is strongly NP-hard when K is part of the input. ■

The mixed-integer solver and the margin test above jointly implement the deterministic linear oracle in (SEP-$\mathscr{H}^{\mathrm{Lin}}$).

H.2 CES, Cobb-Douglas, and Leontief cases: I, Hardness results

We consider $\sigma \in [-1, D_\sigma]$ and impose the parameter box

$$\|\mathbf{y}\|_\infty \leq D_{\mathbf{y}}. \tag{119}$$

The computational results below concern this box-constrained oracle. Every point in the box satisfies $\|\mathbf{y}\|_2 \leq \sqrt{n} D_{\mathbf{y}}$, so Assumption D.1 applies with Euclidean parameter bound $\sqrt{n} D_{\mathbf{y}}$. The Leontief and Cobb-Douglas classes are the special cases. The SEP can be compactly written as follows:

$$\pi(\mathbf{y}, \sigma) := \sum_{k=1}^{K} \langle \mathbf{V}\mathbf{e}_k, \mathrm{softmax}\,(\mathbf{y} - \sigma \log \mathbf{p}_k)\rangle. \tag{120}$$

Now, we first consider the case when σ is fixed:

$$\pi^*(\sigma) = \max_{\|\mathbf{y}\|_\infty \leq D_{\mathbf{y}}} \pi(\mathbf{y}, \sigma),$$

which separates into:

(1) If $\sigma = 0$, then we consider the Cobb-Douglas class, and it becomes,

$$\pi^* = \max_{\|\mathbf{y}\|_\infty \leq D_{\mathbf{y}}} \langle \sum_{k=1}^{K} \mathbf{V}\mathbf{e}_k, \mathrm{softmax}\,(\mathbf{y})\rangle;$$

there is no difference between $K = 1$ and $K \geq 1$ and is always globally solvable in polynomial time.

(2) Otherwise, it is equivalent to solving the following Linear Fractional Programming (LFP) problem with nonnegative coefficients:

$$\vartheta^* = \max_{\mathbf{q} \in \mathcal{Q}} \vartheta(\mathbf{q}) = \sum_{k=1}^{K} \frac{\langle \mathbf{q}, \boldsymbol{\theta}_k \rangle}{\langle \mathbf{q}, \mathbf{f}_k \rangle}, \tag{LFP}$$

where $\boldsymbol{\theta}_k \in \mathbb{R}^n_+$, $\mathbf{f}_k \in \mathbb{R}^n_+$ and $\mathcal{Q}$ is a compact set in $\mathbb{R}^n$ specified in the following result.

Lemma H.1 (SEP with fixed σ is equivalent to LFP)**.** *For any fixed $\sigma \in [-1, D_\sigma]$ and dual iterate $\mathbf{V} \in \mathbb{R}^{n \times K}$, define the following quantities:*

$$\mathbf{q} = \exp(\mathbf{y}) \in \mathbb{R}^n_{>0}, \quad \mathbf{f}_k = \mathbf{p}_k^{-\sigma} \in \mathbb{R}^n_+, \quad \boldsymbol{\theta}_k = \mathrm{diag}(\mathbf{f}_k)\mathbf{V}\mathbf{e}_k \in \mathbb{R}^n_+,$$

and $\mathcal{Q} = \left[e^{-D_{\mathbf{y}}}, e^{D_{\mathbf{y}}}\right]^n \subseteq \mathbb{R}^n_{>0}$. Then, it holds that, $\pi^*(\sigma) = \vartheta^*$.

Proof. Substituting $\boldsymbol{\gamma}(\mathbf{p}_k) = \operatorname{softmax}(\mathbf{y} - \sigma \log \mathbf{p}_k)$ and $\mathbf{q} = \exp(\mathbf{y})$, each summand satisfies

$$\langle \mathbf{V}\mathbf{e}_k, \boldsymbol{\gamma}(\mathbf{p}_k)\rangle = \frac{\langle \mathbf{V}\mathbf{e}_k, \exp(\mathbf{y} - \sigma \log \mathbf{p}_k)\rangle}{\langle \mathbf{1}, \exp(\mathbf{y} - \sigma \log \mathbf{p}_k)\rangle} = \frac{\langle \mathbf{V}\mathbf{e}_k, \operatorname{diag}(\mathbf{p}_k^{-\sigma})\exp(\mathbf{y})\rangle}{\langle \mathbf{1}, \operatorname{diag}(\mathbf{p}_k^{-\sigma})\exp(\mathbf{y})\rangle}$$
$$= \frac{\langle \mathbf{q}, \operatorname{diag}(\mathbf{f}_k)\mathbf{V}\mathbf{e}_k\rangle}{\langle \mathbf{q}, \mathbf{f}_k\rangle} = \frac{\langle \mathbf{q}, \boldsymbol{\theta}_k\rangle}{\langle \mathbf{q}, \mathbf{f}_k\rangle},$$

giving the desired reformulation. In view of Fact 7, $\mathbf{f}_k \in \mathbb{R}^n_+$, and thus $\boldsymbol{\theta}_k \in \mathbb{R}^n_+$. ■

Lemma H.1 reformulates SEP as a sum of linear ratios over its exact feasible box. Hardness of a sum-of-ratios problem over a general polytope (Matsui, 1996; Freund & Jarre, 2001) does not establish hardness of this restricted problem. We instead reduce unweighted Max-Cut directly to the box-constrained SEP, with strictly positive normalized prices and feasible oracle weights.

Theorem H.2 (Complexity of the box-constrained SEP)**.** *Consider SEP on (119) at fixed $\sigma \in [-1, D_\sigma]$. If $K = 1$ or $\sigma = 0$, it reduces to a linear program. For every fixed nonzero rational σ, it is NP-hard when K is part of the input, even with $D_{\mathbf{y}} = \log 2$, positive rational prices $\mathbf{p}_k \in \operatorname{int}\Delta_n$, and $\mathbf{V}\mathbf{e}_k$ a canonical basis vector.*

Proof. If $K = 1$, Lemma H.1 gives one linear-fractional objective on a box, to which the Charnes-Cooper transformation applies (Charnes & Cooper, 1962). If $\sigma = 0$, all price factors equal one, so the K numerators combine into one linear-fractional objective with denominator $\langle \mathbf{1}, \mathbf{q}\rangle$.

For hardness, take an unweighted Max-Cut instance $G = ([n], E)$ with $m = |E| \geq 1$ and target cut size B (Garey et al., 1976). Set

$$D_{\mathbf{y}} := \log 2, \qquad K := 2m, \qquad \Delta := \frac{27}{170}. \tag{121}$$

Thus $\mathbf{q} = \exp(\mathbf{y})$ ranges over exactly $[\frac{1}{2}, 2]^n$. First consider the ideal objective

$$\Phi_0(\mathbf{y}) := \sum_{\{i,j\}\in E} \left(\frac{q_i}{q_i + 4q_j} + \frac{q_j}{q_j + 4q_i} \right). \tag{122}$$

Each ratio is a logistic function of an affine expression in $\mathbf{y}$:

$$\psi(s) := \frac{e^s}{1 + e^s},$$
$$\frac{q_i}{q_i + 4q_j} = \psi(y_i - y_j - \log 4),$$
$$\psi''(s) = \frac{e^s(1 - e^s)}{(1 + e^s)^3} \geq 0, \qquad s \leq 0. \tag{123}$$

Since $y_i - y_j \leq 2D_{\mathbf{y}} = \log 4$, the ideal objective is convex on the entire parameter box and attains a maximum at a box vertex. At such a vertex, an uncut edge contributes $2/5$, while a cut edge contributes $1/2 + 1/17 = 19/34$. Consequently,

$$\max_{\|\mathbf{y}\|_\infty \leq \log 2} \Phi_0(\mathbf{y}) = \frac{2m}{5} + \left(\frac{19}{34} - \frac{2}{5}\right) \operatorname{MaxCut}(G)$$
$$= \frac{2m}{5} + \Delta \operatorname{MaxCut}(G). \tag{124}$$

To obtain strictly positive denominator coefficients, create the two oriented samples (i, j) and (j, i) for each edge and set

$$\tau := \frac{1}{100nm},$$
$$\tilde{f}_{k,i} := 1, \qquad \tilde{f}_{k,j} := 4, \qquad \tilde{f}_{k,\ell} := \tau \quad (\ell \neq i, j),$$
$$\mathbf{V}\mathbf{e}_k := \mathbf{e}_i, \qquad k = (i, j). \tag{125}$$

Let the perturbed objective be

$$\Phi_\tau(\mathbf{y}) := \sum_{k=(i,j)} \frac{q_i}{\langle \tilde{\mathbf{f}}_k, \mathbf{q}\rangle}. \tag{126}$$

For every point in the same box, the loss in one summand satisfies

$$\begin{aligned}\frac{q_i}{q_i+4q_j}-\frac{q_i}{\langle\tilde{\mathbf{f}}_k,\mathbf{q}\rangle}&=\frac{q_i\tau\sum_{\ell\neq i,j}q_\ell}{(q_i+4q_j)\langle\tilde{\mathbf{f}}_k,\mathbf{q}\rangle}\\&\le\frac{\tau\sum_{\ell\neq i,j}q_\ell}{q_i+4q_j}\\&\le\frac{2n\tau}{5/2}\le n\tau.\end{aligned}\tag{127}$$

Summing over the $2m$ samples yields

$$0\le\Phi_0(\mathbf{y})-\Phi_\tau(\mathbf{y})\le 2mn\tau=\frac{1}{50}.\tag{128}$$

It remains to realize these ratios using rational prices for the prescribed fixed nonzero rational σ. Define

$$\rho:=\frac{1}{1000K}.\tag{129}$$

The positive numbers $\tilde{f}_{k,\ell}^{-1/\sigma}$ are algebraic of fixed degree and polynomial encoding length. By bisection to polynomially many bits, choose positive rational numbers $\hat{z}_{k,\ell}$ satisfying

$$(1-\rho)\tilde{f}_{k,\ell}\le\hat{z}_{k,\ell}^{-\sigma}\le(1+\rho)\tilde{f}_{k,\ell},\qquad\ell\in[n],\tag{130}$$

and normalize them into prices:

$$\mathbf{p}_k:=\frac{\hat{\mathbf{z}}_k}{\langle\mathbf{1},\hat{\mathbf{z}}_k\rangle}\in\operatorname{int}\Delta_n.\tag{131}$$

The common normalization factor cancels in each expenditure share, so the actual SEP summand is

$$\begin{aligned}\langle\mathbf{V}\mathbf{e}_k,\operatorname{softmax}(\mathbf{y}-\sigma\log\mathbf{p}_k)\rangle&=\frac{q_ip_{k,i}^{-\sigma}}{\sum_\ell q_\ell p_{k,\ell}^{-\sigma}}\\&=\frac{q_i\hat{z}_{k,i}^{-\sigma}}{\sum_\ell q_\ell\hat{z}_{k,\ell}^{-\sigma}}.\end{aligned}\tag{132}$$

By (130), its numerator and denominator differ from those of $q_i/\langle\tilde{\mathbf{f}}_k,\mathbf{q}\rangle$ by multiplicative factors in $[1-\rho,1+\rho]$. Since this latter ratio lies in $[0,1]$, the absolute change is at most $2\rho/(1-\rho)\le 4\rho$. Therefore, uniformly over the entire feasible box,

$$\begin{aligned}|\pi(\mathbf{y},\sigma)-\Phi_0(\mathbf{y})|&\le|\pi(\mathbf{y},\sigma)-\Phi_\tau(\mathbf{y})|+|\Phi_\tau(\mathbf{y})-\Phi_0(\mathbf{y})|\\&\le 4K\rho+\frac{1}{50}\\&=\frac{3}{125}<\frac{\Delta}{4}.\end{aligned}\tag{133}$$

The rational decision threshold

$$b_G:=\frac{2m}{5}+\left(B-\frac{1}{2}\right)\Delta\tag{134}$$

preserves the answer to Max-Cut:

$$\begin{aligned}\operatorname{MaxCut}(G)\ge B&\implies\pi^*(\sigma)\ge\frac{2m}{5}+B\Delta-\frac{3}{125}>b_G,\\\operatorname{MaxCut}(G)\le B-1&\implies\pi^*(\sigma)\le\frac{2m}{5}+(B-1)\Delta+\frac{3}{125}<b_G.\end{aligned}\tag{135}$$

All prices and oracle weights have polynomial encoding length; the prices are bounded below by an inverse polynomial for fixed σ, so a polynomial bound $D_{\mathbf{p}}$ on $\|\log\mathbf{p}_k\|_2$ suffices. Moreover, $\|\mathbf{V}\mathbf{e}_k\|_1=1$ and the numerator coefficients are exactly $\operatorname{diag}(\mathbf{p}_k^{-\sigma})\mathbf{V}\mathbf{e}_k$. The reduction uses the prescribed box throughout and runs in polynomial time. ■

This result concerns K growing with the input and is compatible with the fixed-K approximation guarantee in Theorem H.3.

H.3 CES, Cobb-Douglas, and Leontief cases: II, An FPTAS for global maximum

Based on the "AppLFMP" algorithm in Depetrini & Locatelli (2011), we can construct a FPTAS to the joint problem in $(\mathbf{y}, \sigma)$ as shown in Algorithm 3. The idea is simple. By constructing a proper

Algorithm 3: Grid-Search FPTAS for SEP

Input: Dual iterate $\mathbf{V}$; tolerance $\epsilon \in (0,1)$.

Set $h \leftarrow \frac{\epsilon}{KD_{\mathbf{p}}}$ and the set of grid points:

$$\Sigma_h \leftarrow \{-1, -1+h, -1+2h, \ldots\} \cap [-1, D_\sigma].$$

foreach $\sigma' \in \Sigma_h$ **do**

  Solve LFP at parameter σ' via "AppLFMP" at relative tolerance $\delta \leftarrow \frac{\epsilon}{2K}$, obtaining $\tilde{\mathbf{q}}(\sigma') \in \mathcal{Q}$;

  Set $\tilde{\mathbf{y}}(\sigma') \leftarrow \log \tilde{\mathbf{q}}(\sigma')$;

end

Set:

$$\hat{\sigma} \leftarrow \arg\max_{\sigma' \in \Sigma_h} \pi(\tilde{\mathbf{y}}(\sigma'), \sigma'), \quad \hat{\mathbf{y}} \leftarrow \tilde{\mathbf{y}}(\hat{\sigma}).$$

Output: $(\hat{\mathbf{y}}, \hat{\sigma})$.

one-dimensional grid of σ, we solve the pricing problem for a fixed σ at each each of the grid point to the relative tolerance $\frac{\epsilon}{2K}$. Then, we aggregate the results to get the final solution by picking the best grid point. We show that Algorithm 3 is an FPTAS.

Theorem H.3. *For any $\epsilon \in (0,1)$ and any fixed K, Algorithm 3 returns an approximate solution $(\hat{\mathbf{y}}, \hat{\sigma})$ in $\tilde{\mathcal{O}}(\epsilon^{-(K+1)})$ time[5] such that $\pi(\hat{\mathbf{y}}, \hat{\sigma}) \geq \pi^* - \epsilon$.*

The first step is to show the complexity for a fixed σ.

H.3.1 An FPTAS for SEP at fixed σ.

Note that Depetrini & Locatelli (2011) considers the linear fractional minimization problem (we are solving a maximization problem). They provided a FPTAS, called "**AppLFMP**", that returns a "relatively" approximate optimal solution if the problem data is integral. To use AppLFMP here, we need to justify that the methods finds an approximate optimal solution measured in *absolute error* for *real-valued data*. We first present the result to guarantee absolute error if the problem data is integral, matching the integral data used in Depetrini & Locatelli (2011).

Lemma H.2. *For any $\sigma \in [-1, D_\sigma]$ and any $\eta \in (0,1)$, the FPTAS of Depetrini & Locatelli (2011) applied to LFP at relative tolerance $\delta = \frac{\eta}{K}$ returns $\tilde{\mathbf{y}} \in [-D_{\mathbf{y}}, D_{\mathbf{y}}]^n$ with*

$$\pi(\tilde{\mathbf{y}}, \sigma) \geq \pi^*(\sigma) - \eta, \tag{136}$$

in $\tilde{\mathcal{O}}(\eta^{-K})$ arithmetic operations.

Proof. First, we justify that the optimal value of LFP is bounded from above. Since $\boldsymbol{\gamma}_k \in \Delta_n$ and $\|\mathbf{V}\mathbf{e}_k\|_2 \leq 1$, each summand of $\pi(\cdot, \sigma)$ is bounded by 1, so the optimal value of LFP is bounded by the sample size.

$$\vartheta^* = \pi^*(\sigma) \leq K. \tag{137}$$

That is, solving LFP is equivalent to solving

$$\min_{\mathbf{q} \in \mathcal{Q}} \; K - \vartheta(\mathbf{q}). \tag{138}$$

Depetrini & Locatelli (2011) considers a δ-relatively optimal solution $\tilde{\mathbf{q}}$ for the minimization problem, which to be precise, refers to $\tilde{\mathbf{q}}$ satisfying the following inequality when applying to (138):

$$(K - \vartheta(\tilde{\mathbf{q}})) \leq (1+\delta) \cdot (K - \vartheta^*),$$

which is equivalent to say,

$$\vartheta(\tilde{\mathbf{q}}) + \delta K \geq (1+\delta) \cdot \vartheta^*. \tag{139}$$

[5] We use $\tilde{\mathcal{O}}(\cdot)$ to hide polynomial terms in fixed parameters.

It is shown that AppLFMP outputs a solution $\tilde{\mathbf{q}}$ satisfying (139) in $\tilde{\mathcal{O}}(\delta^{-K})$ arithmetic operations. By setting $\delta = \frac{\eta}{K}$, AppLFMP returns some $\tilde{\mathbf{q}} \in \mathcal{Q}$ satisfying the following inequality:

$$\vartheta(\tilde{\mathbf{q}}) \geq (1+\delta)\vartheta^* - \delta K \geq \vartheta^* - \delta K = \vartheta^* - \eta,$$

where the second inequality uses $\delta\vartheta^* \geq 0$. By Lemma H.1, we recover the parameter via $\tilde{\mathbf{y}} = \log \tilde{\mathbf{q}}$, which gives the stated bound. The runtime is $\tilde{\mathcal{O}}(\delta^{-K}) = \tilde{\mathcal{O}}(\eta^{-K})$. ■

Next, we show how to apply "AppLFMP" to real-valued coefficients $(\mathbf{f}_k, \boldsymbol{\theta}_k)$. The idea is to find a rounded instance $(\hat{\mathbf{f}}_k, \hat{\boldsymbol{\theta}}_k)$ with bounded encoding length and apply the method to it. Namely, we choose some $\nu \in (0,1), \eta_1 \in (0,1)$,

$$\nu := \frac{\eta_1}{8K} \exp\left(-2D_{\mathbf{y}} - D_\sigma D_{\mathbf{p}}\right) \leq \eta_1; \tag{140}$$

and based on that, set coefficients and objective as follows:

$$\begin{aligned} &\hat{\mathbf{f}}_k := \left\lfloor \tfrac{1}{\nu}\mathbf{f}_k + \tfrac{1}{2}\mathbf{1} \right\rfloor \in \mathbb{N}^n, \qquad \hat{\boldsymbol{\theta}}_k := \left\lfloor \tfrac{1}{\nu}\boldsymbol{\theta}_k + \tfrac{1}{2}\mathbf{1} \right\rfloor \in \mathbb{N}^n, \qquad \forall k \in [K] \\ &\hat{\vartheta}(\mathbf{q}) := \sum_{k=1}^{K} \frac{\langle \mathbf{q}, \hat{\boldsymbol{\theta}}_k \rangle}{\langle \mathbf{q}, \hat{\mathbf{f}}_k \rangle}. \end{aligned} \tag{141}$$

It is legitimate to do so since the ratio computed from $(\hat{\mathbf{f}}_k, \hat{\boldsymbol{\theta}}_k)$ produces the same value as the one from $(\nu\hat{\mathbf{f}}_k, \nu\hat{\boldsymbol{\theta}}_k)$. Using the choice (140), we show that $(\nu\hat{\mathbf{f}}_k, \nu\hat{\boldsymbol{\theta}}_k)$ is close to the original data $(\mathbf{f}_k, \boldsymbol{\theta}_k)$, and the error of a rounding instance can be bounded as follows.

Lemma H.3. *Under Assumption D.1, consider the rounded instance* (141)*, then it satisfies, for all* $\mathbf{q} \in \mathcal{Q}$ *and* $k \in [K]$*,*

$$\|\nu\hat{\mathbf{f}}_k - \mathbf{f}_k\|_\infty \leq \tfrac{\nu}{2}, \quad \|\nu\hat{\boldsymbol{\theta}}_k - \boldsymbol{\theta}_k\|_\infty \leq \tfrac{\nu}{2}, \quad |\vartheta(\mathbf{q}) - \hat{\vartheta}(\mathbf{q})| \leq \tfrac{\eta_1}{2}. \tag{142}$$

Each entry of $\hat{\mathbf{f}}_k, \hat{\boldsymbol{\theta}}_k$ *has encoding length in at most* $\mathcal{O}(\log(\frac{K}{\eta_1}) + D_{\mathbf{y}} + D_\sigma D_{\mathbf{p}})$.

Proof. Under Assumption D.1, and $\mathbf{Ve}_k \in [0,1]^n$, we have

$$\|\log \mathbf{p}_k\|_\infty \leq \|\log \mathbf{p}_k\|_2 \leq D_{\mathbf{p}}, \quad |\sigma| \leq D_\sigma, \quad \mathcal{Q} = [\exp(-D_{\mathbf{y}}), \exp(D_{\mathbf{y}})]^n.$$

Thus, it holds that

$$\exp(-D_\sigma D_{\mathbf{p}})\mathbf{1} \leq \mathbf{f}_k \leq \exp(D_\sigma D_{\mathbf{p}})\mathbf{1}, \qquad \mathbf{0} \leq \boldsymbol{\theta}_k \leq \mathbf{f}_k. \tag{143}$$

By definition (141), $\hat{f}_{k,j} = \left\lfloor \frac{f_{k,j}}{\nu} + \frac{1}{2} \right\rfloor$, so

$$\left| \hat{f}_{k,j} - \tfrac{f_{k,j}}{\nu} \right| \leq \tfrac{1}{2} \quad \Longrightarrow \quad |\nu\hat{f}_{k,j} - f_{k,j}| \leq \tfrac{\nu}{2}, \qquad \forall j \in [n]. \tag{144}$$

Taking the max over j gives $\|\nu\hat{\mathbf{f}}_k - \mathbf{f}_k\|_\infty \leq \frac{\nu}{2}$. $\|\nu\hat{\boldsymbol{\theta}}_k - \boldsymbol{\theta}_k\|_\infty \leq \frac{\nu}{2}$ follows from a similar argument. Since $\mathbf{q} \in \mathcal{Q}$, take

$$A_k := \langle \mathbf{q}, \boldsymbol{\theta}_k \rangle, \quad B_k := \langle \mathbf{q}, \mathbf{f}_k \rangle, \quad \hat{A}_k := \langle \mathbf{q}, \hat{\boldsymbol{\theta}}_k \rangle, \quad \hat{B}_k := \langle \mathbf{q}, \hat{\mathbf{f}}_k \rangle.$$

Then by (143) and (144),

$$B_k \geq n \exp(-D_{\mathbf{y}} - D_\sigma D_{\mathbf{p}}), \quad A_k \leq B_k, \tag{145a}$$

$$|A_k - \nu\hat{A}_k| \leq \|\mathbf{q}\|_1 \tfrac{\nu}{2} \leq \tfrac{n}{2}\exp(D_{\mathbf{y}})\nu \tag{145b}$$

$$|B_k - \nu\hat{B}_k| \leq \|\mathbf{q}\|_1 \tfrac{\nu}{2} \leq \tfrac{n}{2}\exp(D_{\mathbf{y}})\nu. \tag{145c}$$

By the choice (140), we note,

$$\tfrac{n}{2}\exp(D_{\mathbf{y}})\nu = \tfrac{n}{2}\exp(D_{\mathbf{y}})\tfrac{\eta_1}{8K}\exp\left(-2D_{\mathbf{y}} - D_\sigma D_{\mathbf{p}}\right) \leq \tfrac{B_k}{2}.$$

Hence,

$$\begin{aligned} \nu\hat{B}_k &\underset{(145c)}{\geq} B_k - \tfrac{n}{2}\exp(D_{\mathbf{y}})\nu \\ &\geq \tfrac{B_k}{2} \geq \tfrac{n}{2}\exp(-D_{\mathbf{y}} - D_\sigma D_{\mathbf{p}}). \end{aligned} \tag{146}$$

Therefore

$$\left|\frac{A_k}{B_k}-\frac{\nu\hat{A}_k}{\nu\hat{B}_k}\right| = \left|\frac{A_k\cdot\nu\hat{B}_k - B_k\cdot\nu\hat{A}_k}{\nu\hat{B}_k\cdot B_k}\right| = \left|\frac{A_k(\nu\hat{B}_k-B_k)-B_k(\nu\hat{A}_k-A_k)}{\nu\hat{B}_k\cdot B_k}\right|$$
$$\le \frac{A_k}{\nu\hat{B}_k\cdot B_k}|B_k-\nu\hat{B}_k| + \frac{1}{\nu\hat{B}_k}|A_k-\nu\hat{A}_k|$$
$$\underset{(145)}{\le} \left(\frac{A_k}{B_k}+1\right)\cdot\frac{1}{\nu\hat{B}_k}\cdot\frac{n}{2}\exp(D_{\mathbf{y}})\nu \le \frac{n}{\nu\hat{B}_k}\cdot\exp(D_{\mathbf{y}})\nu$$
$$\underset{(146)}{\le} \frac{n\exp(D_{\mathbf{y}})\nu}{\frac{n}{2}\exp(-D_{\mathbf{y}}-D_\sigma D_{\mathbf{p}})} \le 2\exp\left(2D_{\mathbf{y}}+D_\sigma D_{\mathbf{p}}\right)\nu$$
$$= \frac{\eta_1}{4K}. \tag{147}$$

Summing over $k\in[K]$ yields (142). For the encoding length, each entry of $\hat{\mathbf{f}}_k, \hat{\boldsymbol{\theta}}_k$ is a non-negative integer bounded by $\nu^{-1}\exp(D_\sigma D_{\mathbf{p}})$, so its binary encoding takes

$$\text{(bit-length)} = \mathcal{O}\left(\log_2(\nu^{-1}\exp(D_\sigma D_{\mathbf{p}}))\right)$$
$$= \mathcal{O}\left(\log_2\nu^{-1} + D_\sigma D_{\mathbf{p}}\right)$$
$$\underset{(140)}{=} \mathcal{O}\left(\log_2\left(\frac{8K}{\eta_1}\exp\left(2D_{\mathbf{y}}+D_\sigma D_{\mathbf{p}}\right)\right)+D_\sigma D_{\mathbf{p}}\right)$$
$$= \mathcal{O}\left(\log(\tfrac{K}{\eta_1})+D_{\mathbf{y}}+D_\sigma D_{\mathbf{p}}\right). \tag{148}$$

This completes the proof. ■

The complexity for fixed σ is apparent from above results.

Corollary H.1 ("AppLFMP" on real-valued data)**.** *Under Assumption D.1, applying "AppLFMP" to the rounded instance (141) with $\eta_1=\frac{2\eta}{3}$. Set relative tolerance $\delta=\frac{\eta}{3K}$, then it returns $\tilde{\mathbf{y}}\in[-D_{\mathbf{y}},D_{\mathbf{y}}]^n$ in $\tilde{\mathcal{O}}(\eta^{-K})$ time such that*

$$\pi(\tilde{\mathbf{y}},\sigma)\ge\pi^*(\sigma)-\eta. \tag{149}$$

Proof. By Lemma H.3,

$$|\vartheta(\mathbf{q})-\hat{\vartheta}(\mathbf{q})|\le\tfrac{\eta_1}{2}=\tfrac{\eta}{3}.$$

By Lemma H.2 applied to the rounded instance, "AppLFMP" returns $\tilde{\mathbf{q}}\in\mathcal{Q}$ such that

$$\hat{\vartheta}(\tilde{\mathbf{q}})\ge\hat{\vartheta}^*-\tfrac{\eta}{3}.$$

Combining,

$$\vartheta(\tilde{\mathbf{q}})\underset{(142)}{\ge}\hat{\vartheta}(\tilde{\mathbf{q}})-\tfrac{\eta}{3}\underset{\text{Lemma H.2}}{\ge}\hat{\vartheta}^*-\tfrac{2\eta}{3}\underset{(142)}{\ge}\vartheta^*-\eta.$$

Setting $\tilde{\mathbf{y}}:=\log\tilde{\mathbf{q}}$ and using $\pi(\mathbf{y},\sigma)=\vartheta(\exp(\mathbf{y}))$ from Lemma H.1 gives (149). The runtime follow from Lemma H.3. ■

H.3.2 Proof of Theorem H.3.

Finally, we are ready to show the worst-case complexity of Algorithm 3.

Proof. Recall that $\boldsymbol{\gamma}_k(\mathbf{y},\sigma):=\mathrm{softmax}\left(\mathbf{y}-\sigma\log\mathbf{p}_k\right)$, and

$$\pi(\mathbf{y},\sigma):=\sum_{k=1}^{K}\langle\mathbf{V}\mathbf{e}_k,\boldsymbol{\gamma}_k(\mathbf{y},\sigma)\rangle,\quad \pi^*(\sigma):=\max_{\mathbf{y}}\pi(\mathbf{y},\sigma).$$

By Lemma D.1, the softmax is $\frac{1}{2}$-Lipschitz continuous, so for any fixed $\mathbf{y}$ and any σ_1,σ_2,

$$\|\boldsymbol{\gamma}_k(\mathbf{y},\sigma_1)-\boldsymbol{\gamma}_k(\mathbf{y},\sigma_2)\|_2\le\tfrac{1}{2}\|\log\mathbf{p}_k\|_2|\sigma_1-\sigma_2|\le\tfrac{D_{\mathbf{p}}}{2}|\sigma_1-\sigma_2|, \tag{150}$$

Combined with the dual-norm constraint $\|\mathbf{V}\mathbf{e}_k\|_* \le 1$,

$$|\pi(\mathbf{y},\sigma_1)-\pi(\mathbf{y},\sigma_2)| \le \sum_{k=1}^{K} \|\mathbf{V}\mathbf{e}_k\|_2 \|\boldsymbol{\gamma}_k(\mathbf{y},\sigma_1)-\boldsymbol{\gamma}_k(\mathbf{y},\sigma_2)\|_2 \le \tfrac{KD_{\mathbf{P}}}{2}|\sigma_1-\sigma_2|. \tag{151}$$

Since the bound is uniform in $\mathbf{y}$, picking $\mathbf{y}_i^* \in \arg\max_{\mathbf{y}} \pi(\mathbf{y},\sigma_i)$ for $i=1,2$ yields

$$\pi^*(\sigma_1)-\pi^*(\sigma_2) = \pi(\mathbf{y}_1^*,\sigma_1)-\pi(\mathbf{y}_2^*,\sigma_2) \le \pi(\mathbf{y}_1^*,\sigma_1)-\pi(\mathbf{y}_1^*,\sigma_2) \le \tfrac{KD_{\mathbf{P}}}{2}|\sigma_1-\sigma_2|,$$
$$\pi^*(\sigma_2)-\pi^*(\sigma_1) = \pi(\mathbf{y}_2^*,\sigma_2)-\pi(\mathbf{y}_1^*,\sigma_1) \le \pi(\mathbf{y}_2^*,\sigma_2)-\pi(\mathbf{y}_2^*,\sigma_1) \le \tfrac{KD_{\mathbf{P}}}{2}|\sigma_1-\sigma_2|,$$

and thus

$$|\pi^*(\sigma_1)-\pi^*(\sigma_2)| \le \tfrac{KD_{\mathbf{P}}}{2}|\sigma_1-\sigma_2|. \tag{152}$$

By definition, in Algorithm 3 with $h = \frac{\epsilon}{KD_{\mathbf{P}}}$, the size of grid is

$$|\Sigma_h| = \mathcal{O}(1+\tfrac{D_\sigma+1}{h}) = \mathcal{O}(1+\tfrac{KD_{\mathbf{P}}D_\sigma}{\epsilon}).$$

For every $\sigma \in [-1, D_\sigma]$, there exists $\sigma' \in \Sigma_h$ with $|\sigma-\sigma'| \le h$, hence by (152)

$$|\pi^*(\sigma)-\pi^*(\sigma')| \le \tfrac{KD_{\mathbf{P}}}{2}h = \tfrac{\epsilon}{2}.$$

Let $\tilde{\mathbf{y}}(\sigma')$ be the output of "AppLFMP" in Algorithm 3. By Corollary H.1, in $\tilde{\mathcal{O}}(\epsilon^{-K})$ operations,

$$\pi(\tilde{\mathbf{y}}(\sigma'),\sigma') \ge \pi^*(\sigma') - \tfrac{\epsilon}{2}. \tag{153}$$

Let $(\mathbf{y}^*,\sigma^*)$ be a joint maximizer. Then, there exists $\sigma'_* \in \Sigma_h$ with $|\sigma^*-\sigma'_*| \le h$; this σ'_* need not equal $\hat{\sigma}$. So we have,

$$\pi(\hat{\mathbf{y}},\hat{\sigma}) \ge \pi(\tilde{\mathbf{y}}(\sigma'_*),\sigma'_*) \underset{(153)}{\ge} \pi^*(\sigma'_*) - \tfrac{\epsilon}{2} \underset{(152)}{\ge} \pi^*(\sigma^*) - \tfrac{KD_{\mathbf{P}}}{2}h - \tfrac{\epsilon}{2}$$
$$= \pi(\mathbf{y}^*,\sigma^*) - \epsilon.$$

The total cost is $|\Sigma_h| \cdot \tilde{\mathcal{O}}(\epsilon^{-K}) = \tilde{\mathcal{O}}(\epsilon^{-(K+1)})$. This completes the proof. ■

H.4 CES, Cobb-Douglas, and Leontief cases: III, Spatial branch-and-bound and KKT points

Global solution by spatial branch-and-bound. The preceding FPTAS has cost $\tilde{\mathcal{O}}(\epsilon^{-(K+1)})$, which makes it impractical when SEP is called repeatedly. We therefore consider a second global approach based on the bounded nonlinear formulation. By the CES representation in Section B.1, we can reformulate (SEP) using sample expenditure shares $\boldsymbol{\gamma}_k$, log shares $\mathbf{z}_k$, and scalar normalizers A_k. For strictly positive prices and a finite parameter box, the joint optimization over $(\mathbf{y},\sigma)$ becomes

$$\begin{aligned}\pi^* = \max_{\mathbf{y},\sigma,\{\boldsymbol{\gamma}_k,\mathbf{z}_k,A_k\}_{k=1}^K} \ & \sum_{k=1}^{K}\langle \mathbf{V}\mathbf{e}_k,\boldsymbol{\gamma}_k\rangle \\ \text{s.t. } & \mathbf{z}_k = \mathbf{y} - \sigma\log(\mathbf{p}_k) - A_k\mathbf{1}, \quad k\in[K], \\ & \boldsymbol{\gamma}_k = \exp(\mathbf{z}_k), \quad \langle \mathbf{1},\boldsymbol{\gamma}_k\rangle = 1, \quad \mathbf{z}_k \le \mathbf{0}, \quad k\in[K], \\ & \underline{y}\mathbf{1} \le \mathbf{y} \le \overline{y}\mathbf{1}, \quad \underline{\sigma} \le \sigma \le D_\sigma. \end{aligned} \tag{154}$$

Here, exponentiation is componentwise, the parameter bounds are finite, $\underline{y} \le \overline{y}$, and $-1 < \underline{\sigma} \le D_\sigma$. The normalization enforces exactly the softmax shares in (21).

Only the exponential equalities are nonconvex. The formulation can be solved by spatial branch-and-bound, which refines linear outer approximations over the bounded variable domains (Gurobi Optimization, LLC, 2026). The parameter box bounds A_k and $\mathbf{z}_k$, while $\mathbf{z}_k \le \mathbf{0}$ keeps all exponential arguments nonpositive. Given a feasible returned pair $(\hat{\mathbf{y}},\hat{\sigma})$ and the solver's global upper bound $\overline{\pi}$, we reevaluate the CES objective and bound its error by

$$\widehat{\pi} := \sum_{k=1}^{K}\langle \mathbf{V}\mathbf{e}_k, \operatorname{softmax}(\hat{\mathbf{y}} - \hat{\sigma}\log(\mathbf{p}_k))\rangle,$$
$$0 \le \pi^* - \widehat{\pi} \le \overline{\pi} - \widehat{\pi}. \tag{155}$$

Up to numerical tolerances, this gap certifies the additive CES oracle error entering E_T in Theorem 4.2, including when a time limit interrupts the solve. The certificate applies to the specified parameter box.

KKT points. The FPTAS is a theoretical global method whose cost grows rapidly with accuracy, while spatial branch-and-bound can leave a wide optimality gap after a substantial time limit. Both methods are too expensive for the repeated SEP calls made by CRM.

For the practical oracle, we solve (SEP) jointly in $(\mathbf{y}, \sigma)$ over a box,

$$\max_{\mathbf{y}\in\mathbb{R}^n, \sigma\in\mathbb{R}} \pi(\mathbf{y},\sigma) := \sum_{k=1}^{K} \langle \mathbf{V}\mathbf{e}_k, \operatorname{softmax}(\mathbf{y} - \sigma \log(\mathbf{p}_k))\rangle,$$
$$\text{s.t. } \mathbf{y} \in [-D_{\mathbf{y}}, D_{\mathbf{y}}]^n, \quad \sigma \in [-1, D_\sigma]. \tag{156}$$

An ϵ-approximate KKT point satisfies

$$\begin{array}{ll} \left|\partial_{y_j}\pi(\mathbf{y},\sigma)\right| \le \epsilon & \text{if } y_j \in (-D_{\mathbf{y}}, D_{\mathbf{y}}), \\ \max\left\{\partial_{y_j}\pi(\mathbf{y},\sigma), 0\right\} \le \epsilon & \text{if } y_j = -D_{\mathbf{y}}, \\ \max\left\{-\partial_{y_j}\pi(\mathbf{y},\sigma), 0\right\} \le \epsilon & \text{if } y_j = D_{\mathbf{y}}, \\ \left|\partial_{\sigma}\pi(\mathbf{y},\sigma)\right| \le \epsilon & \text{if } \sigma \in (-1, D_\sigma), \\ \max\left\{\partial_{\sigma}\pi(\mathbf{y},\sigma), 0\right\} \le \epsilon & \text{if } \sigma = -1, \\ \max\left\{-\partial_{\sigma}\pi(\mathbf{y},\sigma), 0\right\} \le \epsilon & \text{if } \sigma = D_\sigma. \end{array} \tag{157}$$

Remark 7 (Complexity of ϵ-approximate KKT points)**.** *An ϵ-approximate KKT point can be found in $\mathcal{O}(\epsilon^{-1.5})$ time (Bian et al., 2015).*

I Numerical Experiments: Setup and Full Results

In this section, we report the experimental setup and full numerical results. All experiments are conducted on a single machine with an Intel Core i9-14900HX CPU (24 cores / 32 threads), 16GB RAM, and an NVIDIA GeForce RTX 4060 Laptop GPU (8GB), using PyTorch as the backend.

I.1 Setup

The wealth-redistribution problems (WP)–(WD) and the weighted Nash social welfare problem (NSW) are convex optimization problems. We solve the wealth-redistribution and allocation problems using Cardinal Optimizer (COPT) (Ge et al., 2022). For fully-corrective variants, COPT is also used for the master wealth-reweighting problem. All experiments are evaluated with L_∞ loss and seed $s \in \{0, \ldots, 4\}$, and results are reported as mean±standard deviation.

Methods. We implement the Frank-Wolfe method to realize CRM, including,

(1) `FW`: vanilla Frank–Wolfe method, following Jaggi (2013) and as described in Section E.2 (for the smoothing level, see (91)).
(2) `AFW`: away-steps Frank–Wolfe method (Lacoste-Julien & Jaggi, 2015).
(3) `FCFW`: fully-corrective Frank–Wolfe method (see, for example, Jaggi (2013, Algorithm 4))
Since these two methods are well-known, we do not elaborate on it here.
(4) To deal with linear wealth, we name the extensions using a suffix `(AD)`, including `FW(AD)`, `AFW(AD)`, `FCFW(AD)`.
(5) Because the samples may be large, we also implement the mini-batch version of the methods, with a prefix `S-`. This mainly affects the number of samples used in SEP steps. For example, `S-FW` is the mini-batch version of `FW`. By default, the batch size is set to be 50.

For surrogate-market fitting, we restrict the android class to $\mathscr{H}^{\mathrm{CES}}$, which includes the Cobb–Douglas $\mathscr{H}^{\mathrm{Cobbs}}$ and the Leontief $\mathscr{H}^{\mathrm{Leon}}$ subclasses. We provide two solvers for the nonlinear separation problem: GPU-accelerated PyTorch L-BFGS and Gurobi spatial branch-and-bound. Spatial branch-and-bound can take substantially longer to tighten the global upper bound and obtain a certificate, as shown in Table 6. CRM therefore uses L-BFGS for routine SEP calls and switches to Gurobi only when the improvement found by L-BFGS becomes small.

Real market. We generate a real market with n goods and m agents and produce a training set and a test set of revealed preference data. Each of m agents has a CES utility: the parameter r_i is uniformly selected from $[-3.5, 0.8]$, and a coefficient vector $\mathbf{c}_i \in \mathbb{R}^n_+$ uniformly sampled between $[0, 30]^n$ with density 0.1. Hence, every agent's spending share lies in $\mathscr{H}$. The total wealth is 1 and the price is normalized to be in Δ_n. The dataset is denoted as $\boldsymbol{\Xi} = \{(\mathbf{p}_k, \mathbf{g}_k)\}_{k=1}^K$. Unless stated otherwise, the market scale is set to be n=10, m=30, K=300 with constant wealth. The best-performing variant is highlighted in **bold**.

I.2 Ablation Study

I.2.1 Prediction with Frank–Wolfe variants

Now we present the prediction performance of CRM with different Frank-Wolfe varients. Table 3 reports multi-seed results under two step-size regimes for the non-corrective methods: golden-section linesearch, and the short-step schedule η_t=$2/(t$+$2)$ of Jaggi (2013).

Table 3: Frank–Wolfe variants of CRM.

Step	Method	T	Train L_∞	Test L_∞	Iters
Fully-corrective	FCFW	29±13	$(1.03\pm0.85)\times10^{-2}$	$(1.05\pm0.87)\times10^{-2}$	52±19
	SFCFW	80±17	$\mathbf{(4.07\pm4.45)\times10^{-4}}$	$\mathbf{(4.75\pm4.68)\times10^{-4}}$	445±123
Line search	FW	405±375	$(7.87\pm3.11)\times10^{-3}$	$(8.54\pm3.50)\times10^{-3}$	417±373
	AFW	164±122	$(4.91\pm9.98)\times10^{-2}$	$(4.96\pm9.95)\times10^{-2}$	500
	SFW	500±1	$(7.21\pm2.42)\times10^{-3}$	$(7.70\pm2.79)\times10^{-3}$	500
	SAFW	169±93	$(4.98\pm9.89)\times10^{-2}$	$(5.03\pm9.86)\times10^{-2}$	500
Short step $\eta_t = 2/(t+2)$	FW	44±60	$(1.85\pm2.01)\times10^{-1}$	$(1.84\pm2.01)\times10^{-1}$	59±60
	AFW	61±56	$(9.18\pm18.72)\times10^{-2}$	$(9.18\pm18.69)\times10^{-2}$	309±211
	SFW	41±21	$(3.87\pm4.24)\times10^{-2}$	$(3.91\pm4.32)\times10^{-2}$	56±21
	SAFW	78±40	$(4.43\pm8.81)\times10^{-2}$	$(4.49\pm8.76)\times10^{-2}$	486±31

As is shown, except for the fully-corrective variants, line search consistently yields lower train and test risks than the short-step schedule $\eta_t = 2/(t+2)$. Among all the Frank-Wolfe varients, SFCFW achieves the lowest mean train and test risks, while maintaining a relatively small active set. We therefore use SFCFW as the default CRM solver in the following experiments.

I.2.2 Sensitivity to market and sample size

We next vary the number of goods n, the number of agents m, and the training-sample size $K_{\text{train}} \in \{300, 1000\}$. For each market seed, the 300-observation training set is a prefix of the 1,000-observation set, and both are evaluated on the same independent test set of size 300. All cells use the same mini-batch fully corrective method SFCFW. We use batch size B=50 and a 1,800-second limit for K_{train}=300, and scale these to B=150 and 3,600 seconds for K_{train}=1,000.

Table 4: Problem-size ablation for SFCFW under constant wealth.

		$K_{\text{train}} = 300$			$K_{\text{train}} = 1000$		
n	m	T	Train L_∞	Test L_∞	T	Train L_∞	Test L_∞
10	10	48±29	$(1.67\pm1.99)\times10^{-3}$	$(1.90\pm2.01)\times10^{-3}$	70±51	$(4.2\pm4.2)\times10^{-4}$	$(4.4\pm3.8)\times10^{-4}$
10	30	80±17	$(4.07\pm4.45)\times10^{-4}$	$(4.75\pm4.68)\times10^{-4}$	123±15	$(3.8\pm3.0)\times10^{-4}$	$(3.9\pm2.9)\times10^{-4}$
10	50	118±22	$(7.19\pm4.06)\times10^{-4}$	$(9.50\pm5.83)\times10^{-4}$	134±29	$(5.8\pm4.8)\times10^{-4}$	$(6.6\pm6.1)\times10^{-4}$
20	10	61±23	$(2.33\pm1.87)\times10^{-3}$	$(2.68\pm2.25)\times10^{-3}$	65±22	$(2.80\pm1.27)\times10^{-3}$	$(3.01\pm1.23)\times10^{-3}$
20	30	117±17	$(3.19\pm1.06)\times10^{-3}$	$(3.67\pm1.51)\times10^{-3}$	154±37	$(2.14\pm1.07)\times10^{-3}$	$(2.24\pm1.10)\times10^{-3}$
20	50	138±12	$(2.90\pm0.80)\times10^{-3}$	$(3.50\pm1.32)\times10^{-3}$	178±6	$(2.17\pm1.15)\times10^{-3}$	$(2.29\pm1.35)\times10^{-3}$
50	10	102±23	$(1.31\pm0.39)\times10^{-2}$	$(1.57\pm0.36)\times10^{-2}$	121±29	$(8.49\pm2.34)\times10^{-3}$	$(9.32\pm2.47)\times10^{-3}$
50	30	147±23	$(1.01\pm0.39)\times10^{-2}$	$(1.18\pm0.38)\times10^{-2}$	158±40	$(9.94\pm4.12)\times10^{-3}$	$(1.07\pm0.42)\times10^{-2}$
50	50	173±22	$(7.54\pm2.26)\times10^{-3}$	$(9.05\pm2.41)\times10^{-3}$	212±30	$(5.74\pm1.59)\times10^{-3}$	$(6.16\pm1.70)\times10^{-3}$

Increasing K_{train} from 300 to 1000 lowers the mean test error in eight of the nine market configurations. Among these eight configurations, the reduction ranges from approximately 9% to 77%. For $n = 20$ and $m = 10$, the mean test error instead increases by approximately 12%, from 2.68×10^{-3}

to 3.01×10^{-3}, a change that is small relative to the reported standard deviations. Because the larger training set also uses a larger mini-batch and a longer time limit, this comparison reflects the combined effect of more observations and a larger computational budget.

Market dimensionality remains an important source of difficulty: for every fixed m and K_{train}, the mean test error increases with n. In contrast, the dependence on m is not monotone for either training-sample size, so the table does not support a uniform aggregation effect across the tested configurations. The mean active-set size T increases with the number of agents and with the larger training set in every displayed comparison, while remaining below the corresponding number of training observations. Overall, the most consistent patterns are that prediction becomes harder as the number of goods increases and that the learned surrogate uses more atoms as the number of agents or training observations increases.

I.2.3 Comparing L-BFGS to spatial branch-and-bound

Effect of initializations on L-BFGS. Because the CES separation problem is nonconvex, L-BFGS can in principle depend on its initialization. We therefore rerun the same constant-wealth SFCFW experiment with multiple starts per separation call and retain the best feasible separation value. This measures whether initialization changes the final prediction rather than merely improving an individual oracle call.

Table 5: Sensitivity to L-BFGS initialization.

L-BFGS starts	Train L_∞	Test L_∞
1	$(4.07 \pm 4.45) \times 10^{-4}$	$(4.75 \pm 4.68) \times 10^{-4}$
2	$(3.98 \pm 3.04) \times 10^{-4}$	$(4.68 \pm 3.22) \times 10^{-4}$
4	$(3.99 \pm 2.92) \times 10^{-4}$	$(4.89 \pm 3.12) \times 10^{-4}$
8	$(3.87 \pm 3.03) \times 10^{-4}$	$(4.68 \pm 3.35) \times 10^{-4}$

The prediction error is largely insensitive to the number of L-BFGS initializations: all four settings achieve comparable train and test risks. Increasing the number of starts provides little improvement in prediction accuracy but brings additional training time, suggesting that a small number of L-BFGS initializations is sufficient for the separation oracle in this setting.

L-BFGS vs spatial branch-and-bound I: solving separation problems. We run SFCFW and test the first 20 separation instances generated during its iterations, comparing the feasible value returned by L-BFGS with the incumbent and global upper bound returned by Gurobi. We report the mean and maximum relative gaps across all instances. Best feasible values at the displayed precision are shown in bold.

Table 6: L-BFGS values and Gurobi spatial branch-and-bound certificates for 20 CES separation instances. Times, values, and bounds are rounded to four decimal places.

Problem	Method	Time	Value	Bound	Problem	Method	Time	Value	Bound
	L-BFGS	0.0049	**3.5369**			L-BFGS	0.0079	**3.4034**	
1	Gurobi (10s)	3.7605	3.5369	3.5369	2	Gurobi (10s)	0.9603	3.4034	3.4034
	Gurobi (limit)	3.9001	3.5369	3.5369		Gurobi (limit)	0.9371	3.4034	3.4034
	L-BFGS	0.0092	**3.8778**			L-BFGS	0.0073	**5.1069**	
3	Gurobi (10s)	10.0042	3.8750	5.3225	4	Gurobi (10s)	10.0038	5.1069	6.2811
	Gurobi (limit)	127.3380	3.8778	3.8782		Gurobi (limit)	64.4460	5.1069	5.1069
	L-BFGS	0.0072	**4.7890**			L-BFGS	0.0079	3.8971	
5	Gurobi (10s)	10.0039	4.7890	5.8328	6	Gurobi (10s)	7.0664	**4.0126**	4.0129
	Gurobi (limit)	77.4786	4.7890	4.7893		Gurobi (limit)	7.2339	**4.0126**	4.0129
	L-BFGS	0.0091	**3.7084**			L-BFGS	0.0073	**5.3687**	
7	Gurobi (10s)	10.0055	3.7084	4.5315	8	Gurobi (10s)	10.0034	5.3687	6.2891
	Gurobi (limit)	118.8990	3.7083	3.7087		Gurobi (limit)	25.0677	5.3687	5.3690
	L-BFGS	0.0100	**4.0859**			L-BFGS	0.0127	**3.5893**	
9	Gurobi (10s)	10.0198	4.0859	6.4167	10	Gurobi (10s)	10.0051	3.5893	5.1774
	Gurobi (limit)	2000.0200	4.0859	4.0867		Gurobi (limit)	372.0270	3.5893	3.5896
	L-BFGS	0.0067	**3.7460**			L-BFGS	0.0072	**2.5026**	
11	Gurobi (10s)	10.0165	3.7460	5.4927	12	Gurobi (10s)	10.0035	2.5026	3.4748

Continued on next page

Problem	Method	Time	Value	Bound	Problem	Method	Time	Value	Bound
	Gurobi (limit)	562.2210	3.7460	3.7462		Gurobi (limit)	69.5518	2.5026	2.5026
	L-BFGS	0.0069	**3.9239**			L-BFGS	0.0158	**7.0660**	
13	Gurobi (10s)	10.0034	3.9239	4.8698	14	Gurobi (10s)	10.0050	7.0660	8.0712
	Gurobi (limit)	147.7350	3.9239	3.9243		Gurobi (limit)	46.1424	7.0660	7.0667
	L-BFGS	0.0091	**3.1995**			L-BFGS	0.0083	**5.5485**	
15	Gurobi (10s)	10.0037	3.1995	4.8286	16	Gurobi (10s)	10.0045	5.5485	6.2748
	Gurobi (limit)	64.5081	3.1995	3.1998		Gurobi (limit)	34.5877	5.5485	5.5489
	L-BFGS	0.0061	**4.4412**			L-BFGS	0.0121	**3.6990**	
17	Gurobi (10s)	10.0083	4.4412	5.9785	18	Gurobi (10s)	10.0055	3.6990	5.3764
	Gurobi (limit)	87.9424	4.4412	4.4416		Gurobi (limit)	905.3410	3.6990	3.6993
	L-BFGS	0.0105	**3.8637**			L-BFGS	0.0071	**4.5745**	
19	Gurobi (10s)	10.0044	3.8637	4.5212	20	Gurobi (10s)	2.4612	4.5745	4.5745
	Gurobi (limit)	58.9504	3.8637	3.8637		Gurobi (limit)	2.4653	4.5745	4.5745

We report $\frac{|v_{\mathrm{LBFGS}}-v_{\mathrm{inc}}|}{|v_{\mathrm{LBFGS}}|}$ and $\frac{u_{\mathrm{Gurobi}}-v_{\mathrm{LBFGS}}}{|v_{\mathrm{LBFGS}}|}$, where v_{inc} is the best Gurobi incumbent and u_{Gurobi} is the final Gurobi upper bound. The results indicate that L-BFGS is typically close to optimal: it

Table 7: Mean and maximum relative gaps across the 20 CES separation instances.

Comparison	Mean	Maximum
L-BFGS to best Gurobi incumbent	0.15%	2.96%
L-BFGS to final Gurobi upper bound	0.15%	2.97%

matches the best feasible value found by Gurobi to the reported precision on 19 of 20 instances. The only exception is instance 6, where Gurobi improves the feasible value by 2.96% relative to the L-BFGS value. Certifying global optimality is considerably harder because the Gurobi upper bound improves slowly. The largest final relative upper-bound gap is 2.97% on instance 6, where Gurobi also finds the better feasible value, while instance 9 requires 2,000 seconds to reduce its gap below 0.02%.

L-BFGS vs spatial branch-and-bound II: end-to-end comparison in CRM. We compare L-BFGS with Gurobi's spatial branch-and-bound as the separation solver throughout the full CRM-`SFCFW` procedure. Both configurations use constant-wealth CES markets with $n = 10$ goods, $m = 30$ agents, and $K_{\mathrm{train}} = K_{\mathrm{test}} = 300$ observations. Each run uses the L_∞ loss, line search, 300 CRM iterations without patience-based stopping, and the CPU implementation. The solver-specific settings are as follows:

- **L-BFGS:** we test stochastic batch sizes in $\{10, 20, 50\}$ over five market seeds; all rest configurations are the same as before.
- **Gurobi:** we test stochastic batches of 10 observations. We do not test larger batch sizes because the separation problems are computationally difficult and it can take too long to complete. We set time limit to 100 seconds, the absolute optimality-gap target to 10^{-6} (`MIPGapAbs`) and the relative optimality-gap target to 10^{-3} (`MIPGap`), and solve the nonlinear model by spatial branch-and-bound with `NonConvex=2`.

Table 8: End-to-end comparison of separation solvers in CRM-`SFCFW`. Values are the mean and sample standard deviation over five market seeds. At 100-second limit, Gurobi may not reach the gap targets (mean gap cross the iterates is 2.58%.)

Separation solver	Batch	Runs	Train L_∞	Test L_∞	SEP time (s)	Wall time (s)
L-BFGS	10	5	$(8.413\pm6.908)\times10^{-4}$	$(9.037\pm6.722)\times10^{-4}$	10.63±0.92	17.14±1.54
L-BFGS	20	5	$(5.748\pm4.320)\times10^{-4}$	$(6.055\pm4.034)\times10^{-4}$	13.07±1.70	21.06±2.04
L-BFGS	50	5	$(5.940\pm3.792)\times10^{-4}$	$(6.935\pm3.882)\times10^{-4}$	16.24±0.96	28.04±5.26
Gurobi	10	5	$(6.728\pm0.141)\times10^{-4}$	$(7.010\pm0.223)\times10^{-4}$	$>15{,}000$	$>15{,}000$

Even with only 10 observations per batch, the Gurobi separation problems are computationally difficult: the total separation time exceeds 15,000 seconds per run, and the per-call time limit may be reached before either gap target. At the same batch size of 10, Gurobi attains a lower mean test

risk than L-BFGS, 7.010×10^{-4} compared with 9.037×10^{-4}. This result indicates that the global solver obtains a stronger separation direction from the same sampled observations. L-BFGS closes this gap by using batch sizes 20 or 50, which attain test risks of 6.055×10^{-4} and 6.935×10^{-4}, while remaining much faster. The Gurobi separation time is more than 900 times the mean time of every L-BFGS configuration, so we use L-BFGS as the default separation method in the subsequent experiments.

I.2.4 CRM VS OTHER BASELINE METHODS

We compare three ways to learn a share map $\mathbf{p} \mapsto \mathbf{g}(\mathbf{p})$ on constant-wealth CES and GES markets. For each market seed $s \in \{0, \dots, 4\}$, we use $n = 10$ goods and $m = 30$ agents, draw $K_{\text{train}} = 300$ observations with price seed s, and draw an independent $K_{\text{test}} = 300$-point test set with price seed $s + 10000$. Every method is trained with the training L_∞ risk. We report the capacity and parameters of each method, and evaluate both held-out prediction error. For CRM-`SFCFW` and CRM-`SFCFW(AD)`, the reported parameter is $T(n+1)$ where T is the number of active atoms.

Fixed random CES. We draw $D \in \{25, 50, 100\}$ CES atoms from the data-generating distribution using dictionary seed $s + 100000$, freeze their $(\mathbf{y}, \sigma)$ parameters, and fit only the mixture weights. Thus, the reported parameter count is $D(n+1) \in \{275, 550, 1100\}$, excluding mixture weights. This baseline tests a dictionary whose nonlinear features are fixed before fitting.

Simplex-MLP. The network maps $\log p \in \mathbb{R}^n$ through two GELU hidden layers and an n-way softmax. We use full-batch AdamW (learning rate 10^{-3}, weight decay 10^{-5}, with dropout=0.1), at most 2,000 epochs, and training-only early stopping with patience 150. For widths (h_1, h_2), $P = (n+1)h_1 + (h_1+1)h_2 + nh_2 + n$.

Table 9: Comparison of CRM to baseline methods.

Method	Capacity	Parameters	Train L_∞	Test L_∞
CES market				
CRM-`SFCFW`	$T = 80 \pm 17$	880±187	$\mathbf{(4.07 \pm 4.45) \times 10^{-4}}$	$\mathbf{(4.75 \pm 4.68) \times 10^{-4}}$
Fixed random CES	$D = 25$	275	$(4.930 \pm 2.683) \times 10^{-2}$	$(4.956 \pm 2.689) \times 10^{-2}$
Fixed random CES	$D = 50$	550	$(1.850 \pm 0.825) \times 10^{-2}$	$(1.863 \pm 0.876) \times 10^{-2}$
Fixed random CES	$D = 100$	1,100	$(1.490 \pm 0.884) \times 10^{-2}$	$(1.516 \pm 0.926) \times 10^{-2}$
Simplex-MLP	16×16	618	$(4.76 \pm 0.86) \times 10^{-3}$	$(4.94 \pm 1.00) \times 10^{-3}$
Simplex-MLP	32×32	1,738	$(1.70 \pm 0.40) \times 10^{-3}$	$(2.05 \pm 0.54) \times 10^{-3}$
Simplex-MLP	64×64	5,514	$(1.27 \pm 0.23) \times 10^{-3}$	$(1.75 \pm 0.43) \times 10^{-3}$
GES market				
CRM-`SFCFW(AD)`	$T = 126 \pm 23$	1,386±253	$(5.73 \pm 0.47) \times 10^{-3}$	$\mathbf{(2.018 \pm 0.109) \times 10^{-2}}$
Fixed random CES	$D = 25$	275	$(3.418 \pm 0.104) \times 10^{-1}$	$(3.365 \pm 0.092) \times 10^{-1}$
Fixed random CES	$D = 50$	550	$(3.345 \pm 0.145) \times 10^{-1}$	$(3.289 \pm 0.082) \times 10^{-1}$
Fixed random CES	$D = 100$	1,100	$(3.159 \pm 0.179) \times 10^{-1}$	$(3.116 \pm 0.072) \times 10^{-1}$
Simplex-MLP	16×16	618	$(4.39 \pm 0.52) \times 10^{-3}$	$(7.91 \pm 0.66) \times 10^{-2}$
Simplex-MLP	32×32	1,738	$(9.39 \pm 1.01) \times 10^{-4}$	$(4.52 \pm 0.35) \times 10^{-2}$
Simplex-MLP	64×64	5,514	$(2.39 \pm 0.27) \times 10^{-4}$	$(2.832 \pm 0.421) \times 10^{-2}$

As shown, CRM achieves the lowest mean test L_∞ loss across the five seeds in both markets, while the fixed dictionaries substantially underfit. In the CES market, both MLP train and test losses decrease with width, but the best mean MLP test loss is still approximately 3.7 times that of CRM while using approximately 6.3 times as many parameters. In the GES market, increasing the MLP width similarly improves both train and test losses. Although the widest MLP attains a lower training loss than CRM, its mean test loss remains approximately 1.4 times that of CRM while using approximately 4.0 times as many parameters.

I.3 PREDICTION UNDER DIFFERENT WEALTH FUNCTIONS

We now compare the performance of CRM under different wealth functions in the real market. We consider the following wealth functions:

(1) Constant / Zero-order: $w_i(\mathbf{p}) \equiv w_i$ with $\sum_{i\in[m]} w_i = 1$, i.e., the Fisher model.

(2) First-order: $w_i(\mathbf{p}) = \langle \mathbf{p}, \mathbf{b}_i \rangle$ with $\mathbf{b}_i \in \mathbb{R}^n_+$ and $\sum_{i\in[m]} \mathbf{b}_i = \mathbf{1}$, i.e., the Arrow-Debreu model.

(3) Second-order: $w_i(\mathbf{p}) = \frac{\mathbf{p}^\top \mathbf{Q}_i \mathbf{p}}{\sum_{j\in[m]} \mathbf{p}^\top \mathbf{Q}_j \mathbf{p}}$ with symmetric positive semidefinite $\mathbf{Q}_i \succeq 0$; thus we still have $\sum_{i\in[m]} w_i(\mathbf{p}) = 1$, but the homogeneity properties are lost.

Accordingly, the main text reports only `SFCFW` and its linear-wealth extension `SFCFW(AD)`. Under constant real wealth, `SFCFW` attains lower test risk and uses a smaller active set T. Under first- and second-order real wealth, `SFCFW(AD)` attains lower test risk, showing the benefit of allowing the surrogate wealth to depend on prices.

I.4 COMPUTING FAIR ALLOCATIONS

We apply CRM to compute a proportionally fair allocation (PFA) under fixed wealth functions. We use the following simple rule to allocate:

(1) Learn a surrogate demand function from the historical price-aggregate demand pairs $\{(\mathbf{p}_k, \mathbf{d}(\mathbf{p}_k))\}_{k=1}^K$.

(2) Compute the surrogate market's equilibrium price $\mathbf{p}^*$ (cf. WE). This can be done efficiently by a price-adjustment process (Cheung et al., 2020; Zhang et al., 2025) if all the wealth functions are constants in the surrogate market.[6]

(3) Post $\mathbf{p}^*$ and query each current agent i exactly once for its demand $\mathbf{x}_i(\mathbf{p}^*)$.

(4) Rescale the reported demands good by good to satisfy the supply constraints:

$$\widehat{\mathbf{x}}_i = \frac{\mathbf{x}_i(\mathbf{p}^*)}{\sum_{i\in[m]} \mathbf{x}_i(\mathbf{p}^*)}, \quad \widehat{\mathrm{NSW}} := \prod_{i\in[m]} u_i(\widehat{\mathbf{x}}_i)^{w_i}. \tag{158}$$

This protocol uses aggregate data during learning and one round of individual communication during implementation; it does not require full preference elicitation or an iterative exchange with individual agents. By (158), $\sum_{i\in[m]} \widehat{\mathbf{x}}_i = \mathbf{1}$, and so $\{\widehat{\mathbf{x}}_i\}_{i\in[m]}$ is a feasible allocation. For comparison, we compute the optimal allocation $\{\mathbf{x}_i^{\mathrm{opt}}\}_{i\in[m]}$ by the following convex optimization problem:

$$\mathrm{NSW}^{\mathrm{opt}} = \max_{\mathbf{x}_i \in \mathbb{R}^n_+} \prod_{i\in[m]} u_i(\mathbf{x}_i)^{w_i}, \quad \text{s.t.} \sum_{i\in[m]} \mathbf{x}_i \le \mathbf{1}. \tag{159}$$

We must have $\mathrm{NSW}^{\mathrm{opt}} \ge \widehat{\mathrm{NSW}}$. We report the relative loss in Nash social welfare:

$$\Delta_{\mathrm{NSW}} := 1 - \frac{\widehat{\mathrm{NSW}}}{\mathrm{NSW}^{\mathrm{opt}}}. \tag{160}$$

Smaller values mean that the allocation $\{\widehat{\mathbf{x}}_i\}_{i\in[m]}$ is closer in welfare to the PFA. We test two types of markets: (1) CES market, generated in the same way as in Section I.3; (2) non-homothetic market using the following utility function,

$$u_i(\mathbf{x}) = \sum_{j\in[n]} c_{ij}\, x_j^{r_{ij}}, \quad \mathbf{c}_i \in \mathbb{R}^n_+, \quad r_{ij} \in (0,1). \tag{161}$$

The utility can be seen as a generalization of the CES utility by allowing r_{ij} to vary across goods j, making the agents non-homothetic. In CRM, we only use the CES androids and constant wealth. Table 10 reports the results for $n \in \{10, 20, 50, 100\}$ goods and $m \in \{1000, 5000\}$ agents, using $K = 300$ observations and a 1800-second time limit.

[6] While allowing linear wealth makes the problem PPAD-hard (Chen & Teng, 2009; Chen et al., 2017), it is possible to implement a vanilla Newton's method to solve WE as a nonlinear least-squares problem, which is usually capable of obtaining a solution at a precision of 10^{-6} without complexity guarantees. We do not elaborate this option as using constant wealth suffices in our experiments.

Table 10: Fair allocation using SFCFW on constant-wealth CES and non-homothetic GES markets.

Market	m	Metric	$n = 10$	20	50	100
CES	1000	T	175±4	213±12	217±19	183±22
		Δ_{NSW} (%)	$(1.09\pm0.40) \times 10^{-5}$	$(1.25\pm0.34) \times 10^{-3}$	$(4.30\pm1.33) \times 10^{-2}$	$(2.71\pm1.54) \times 10^{-2}$
	5000	T	150±3	186±16	201±6	160±25
		Δ_{NSW} (%)	$(1.28\pm0.59) \times 10^{-5}$	$(7.77\pm2.63) \times 10^{-4}$	$(2.94\pm1.17) \times 10^{-2}$	$(1.39\pm1.02) \times 10^{-2}$
GES	1000	T	230±16	211±4	166±89	158±13
		Δ_{NSW} (%)	$(1.98\pm0.09) \times 10^{-1}$	$(9.58\pm1.22) \times 10^{-2}$	$(1.80\pm2.95) \times 10^{-1}$	$(6.33\pm0.91) \times 10^{-2}$
	5000	T	242±6	221±14	166±6	135±19
		Δ_{NSW} (%)	$(1.80\pm0.06) \times 10^{-1}$	$(8.42\pm0.19) \times 10^{-2}$	$(4.02\pm0.18) \times 10^{-2}$	$(5.79\pm1.71) \times 10^{-2}$

The surrogate uses only CES androids in both settings. All instances have welfare loss below 3%. Furthermore, increasing m does not produce a systematic increase in the number of active androids T. From a computational perspective, this is attractive because both the construction of the surrogate market and the subsequent allocation are driven by aggregate demand rather than individual preferences. By contrast, computing the optimal PFA directly (i.e., by (159)) involves all agents explicitly, whose size scales directly with m. These results therefore suggest that CRM may provide a scalable approach to computing approximately fair allocations when m is large or personal utilities are unavailable.

I.5 Performance on approximate PFA

I.5.1 Setup

Recall that $\mathbf{h}(\mathbf{p})$ is the expenditure share predicted by the surrogate market. In this section, the utilities $u_i : \mathbb{R}^n_+ \to \mathbb{R}_+$ are continuous, monotone, and concave, with $u_i(\mathbf{x}) > 0$ for $\mathbf{x} \in \mathbb{R}^n_{>0}$, and the wealths $w_i > 0$ are fixed with $\sum_i w_i = 1$. The following condition bounds how utility changes when a consumption bundle is scaled.

Assumption I.1 (Bounded elasticity)**.** *There exist $0 < \underline{r} \le \overline{r}$ such that, for every $i \in [m]$ and $\mathbf{x} \ge \mathbf{0}$,*

$$\begin{aligned} t^{\underline{r}} u_i(\mathbf{x}) \le u_i(t\mathbf{x}) \le t^{\overline{r}} u_i(\mathbf{x}), \qquad & t \ge 1, \\ t^{\overline{r}} u_i(\mathbf{x}) \le u_i(t\mathbf{x}) \le t^{\underline{r}} u_i(\mathbf{x}), \qquad & 0 < t \le 1. \end{aligned} \tag{162}$$

When u_i is differentiable on $\mathbb{R}^n_{>0}$, (162) is equivalent to

$$\underline{r} \le \frac{\langle \nabla u_i(\mathbf{x}), \mathbf{x} \rangle}{u_i(\mathbf{x})} \le \overline{r}, \qquad \mathbf{x} \in \mathbb{R}^n_{>0}. \tag{163}$$

Degree-one homogeneous utilities satisfy Assumption I.1 with $\underline{r} = \overline{r} = 1$. For the non-homothetic GES utilities in (161), the condition holds termwise with

$$\underline{r} := \min_{i \in [m], j \in [n]} r_{ij}, \qquad \overline{r} := \max_{i \in [m], j \in [n]} r_{ij}. \tag{164}$$

We assume each GES coefficient vector $\mathbf{c}_i$ is nonzero, so its utility is positive on $\mathbb{R}^n_{>0}$. Define the entropy of the normalized wealths by

$$H(\mathbf{w}) := \sum_{i=1}^m w_i \log \frac{1}{w_i}. \tag{165}$$

The weighted Nash welfare and its feasible optimum are

$$\begin{aligned} \log \mathrm{NSW}(\mathbf{x}) &:= \sum_{i=1}^m w_i \log u_i(\mathbf{x}_i), \\ \log \mathrm{NSW}^{\mathrm{opt}} &:= \max_{\mathbf{x}_i \ge \mathbf{0}} \sum_{i=1}^m w_i \log u_i(\mathbf{x}_i) \\ &\text{s.t.} \sum_{i=1}^m \mathbf{x}_i \le \mathbf{1}. \end{aligned} \tag{166}$$

Fix $\epsilon > 0$, $\delta \in (0,1)$, and $0 < \underline{p} \le \frac{1}{n}$. Define the compact interior price domain

$$\mathcal{P}_{\underline{p}} := \left\{ \mathbf{p} \in \Delta_n : \min_{j\in[n]} p_j \ge \underline{p} \right\}. \tag{167}$$

Assumption I.2 (Learning error, price coverage, and regularity)**.** *The learned surrogate* $\mathbf{h} \in \mathrm{conv}(\mathscr{H})$ *satisfies*

$$L(\mathbf{h}) - L^* \le \epsilon. \tag{168}$$

Moreover, τ *is supported on* $\mathcal{P}_{\underline{p}}$*, and there exist* $c, r_0, G > 0$ *such that*

$$\tau\left(B(\mathbf{p}, r) \cap \mathcal{P}_{\underline{p}}\right) \ge c r^{n-1}, \qquad \mathbf{p} \in \mathcal{P}_{\underline{p}}, \quad 0 < r \le r_0, \tag{169}$$

and the map $\mathbf{p} \mapsto \|\mathbf{g}(\mathbf{p}) - \mathbf{h}(\mathbf{p})\|_\infty$ *is* G*-Lipschitz on* $\mathcal{P}_{\underline{p}}$*.*

This assumption requires the learned surrogate to be empirically near-optimal, the price distribution to cover every local neighborhood of the interior price domain, and the prediction error to vary with price at a controlled rate. Let us define

$$\overline{R} := L^* + \epsilon + 2\sqrt{2}\left(\frac{\mathrm{poly}(n, D_\sigma, D_{\mathbf{y}}, D_{\mathbf{p}})}{\sqrt{K}} + \sqrt{\frac{2\log \bar{f}}{K}}\right) + 6\sqrt{\frac{\ln(2/\delta)}{2K}},$$
$$r := \min\left\{ r_0, \left(\frac{\overline{R}}{cG}\right)^{1/n} \right\}, \qquad \eta := Gr + \frac{\overline{R}}{c r^{n-1}}. \tag{170}$$

By the generalization bound in the main text, with probability at least $1-\delta$,

$$\mathbb{E}_{\mathbf{p}\sim\tau}\left[\|\mathbf{g}(\mathbf{p}) - \mathbf{h}(\mathbf{p})\|_\infty\right] \le \overline{R}. \tag{171}$$

On this event, for every $\mathbf{p} \in \mathcal{P}_{\underline{p}}$,

$$\|\mathbf{g}(\mathbf{p}) - \mathbf{h}(\mathbf{p})\|_\infty \le Gr + \frac{\int_{B(\mathbf{p},r)\cap\mathcal{P}_{\underline{p}}} \|\mathbf{g}(\mathbf{p}') - \mathbf{h}(\mathbf{p}')\|_\infty d\tau(\mathbf{p}')}{\tau(B(\mathbf{p},r)\cap\mathcal{P}_{\underline{p}})}$$
$$\le Gr + \frac{\overline{R}}{c r^{n-1}} = \eta. \tag{172}$$

Assumption I.3 (Equilibrium error)**.** *The computed price* $\mathbf{p}^* \in \Delta_n$ *satisfies*

$$\|\mathbf{h}(\mathbf{p}^*) - \mathbf{p}^*\|_\infty \le \epsilon, \quad \min_{j\in[n]} p_j^* \quad \ge \underline{p}. \tag{173}$$

By Assumption I.3, $\mathbf{p}^* \in \mathcal{P}_{\underline{p}}$. Consequently, although $\mathbf{p}^*$ is computed from $\mathbf{h}$, with probability at least $1-\delta$,

$$\boxed{\|\mathbf{g}(\mathbf{p}^*) - \mathbf{h}(\mathbf{p}^*)\|_\infty \le \eta.} \tag{174}$$

Assume that $\mathbf{d}(\mathbf{p}^*) \in \mathbb{R}^n_{>0}$, so $\mathrm{diag}(\mathbf{d}(\mathbf{p}^*))$ is invertible. Recall that the allocation implemented by CRM is

$$\widehat{\mathbf{x}}_i(\mathbf{p}^*) := [\mathrm{diag}(\mathbf{d}(\mathbf{p}^*))]^{-1}\mathbf{x}_i(\mathbf{p}^*), \qquad \sum_{i=1}^m \widehat{\mathbf{x}}_i(\mathbf{p}^*) = \mathbf{1}. \tag{175}$$

Define the realized prediction and equilibrium errors and the smallest computed price by

$$\eta^* := \|\mathbf{g}(\mathbf{p}^*) - \mathbf{h}(\mathbf{p}^*)\|_\infty, \qquad \xi^* := \|\mathbf{h}(\mathbf{p}^*) - \mathbf{p}^*\|_\infty,$$
$$p^*_{\min} := \min_{j\in[n]} p_j^*. \tag{176}$$

Define the maximum observed aggregate demand and log-welfare gap by

$$M^* := \max_{j\in[n]} d_j(\mathbf{p}^*). \tag{177}$$

$$\Delta_{\log} := \log \frac{\mathrm{NSW}^{\mathrm{opt}}}{\mathrm{NSW}(\widehat{\mathbf{x}}(\mathbf{p}^*))}. \tag{178}$$

I.5.2 MAIN RESULT

Theorem I.1 (Nash-welfare error)**.** *Under Assumption I.1, for any* $\mathbf{p}^* \in \mathbb{R}^n_{>0} \cap \Delta_n$ *with* $\mathbf{d}(\mathbf{p}^*) \in \mathbb{R}^n_{>0}$*, the allocation* (175) *satisfies*

$$0 \le \Delta_{\log} \le (\overline{r} - \underline{r}) \min\{H(\mathbf{w}), 1/e\} + \overline{r} \log M^*,$$
$$0 \le \Delta_{\mathrm{NSW}} \le 1 - e^{-(\overline{r}-\underline{r}) \min\{H(\mathbf{w}),1/e\}} (M^*)^{-\overline{r}}. \tag{179}$$

If Assumptions I.2 and I.3 also hold, then with probability at least $1 - \delta$*,*

$$\begin{aligned} 0 \le \Delta_{\log} &\le (\overline{r} - \underline{r}) \min\{H(\mathbf{w}), 1/e\} + \overline{r} \log \left(1 + \frac{\eta^* + \xi^*}{p^*_{\min}}\right) \\ &\le (\overline{r} - \underline{r}) \min\{H(\mathbf{w}), 1/e\} + \overline{r} \log \left(1 + \frac{\eta + \epsilon}{\underline{p}}\right). \end{aligned} \tag{180}$$

Consequently,

$$0 \le \Delta_{\mathrm{NSW}} \le 1 - e^{-(\overline{r}-\underline{r}) \min\{H(\mathbf{w}),1/e\}} \left(\frac{\underline{p}}{\underline{p} + \eta + \epsilon}\right)^{\overline{r}}, \tag{181}$$

where η *is defined in* (170)*.*

Proof. By the triangle inequality,

$$\begin{aligned} \|\mathrm{diag}(\mathbf{p}^*)(\mathbf{d}(\mathbf{p}^*) - \mathbf{1})\|_\infty &= \|\mathbf{g}(\mathbf{p}^*) - \mathbf{p}^*\|_\infty \\ &\le \|\mathbf{g}(\mathbf{p}^*) - \mathbf{h}(\mathbf{p}^*)\|_\infty + \|\mathbf{h}(\mathbf{p}^*) - \mathbf{p}^*\|_\infty \\ &= \eta^* + \xi^*. \end{aligned} \tag{182}$$

Walras' law and $\sum_{i=1}^m w_i = 1$ imply $1 = \langle \mathbf{p}^*, \mathbf{d}(\mathbf{p}^*)\rangle \le M^* \langle \mathbf{p}^*, \mathbf{1}\rangle = M^*$. By (182),

$$M^* \le 1 + \frac{\eta^* + \xi^*}{p^*_{\min}}. \tag{183}$$

Let $\mathbf{x}^{\mathrm{opt}}$ maximize feasible Nash welfare and define its expenditures at $\mathbf{p}^*$ by

$$b_i := \langle \mathbf{p}^*, \mathbf{x}_i^{\mathrm{opt}} \rangle, \qquad i \in [m]. \tag{184}$$

Positive utility is attainable for every agent under unit supply, so every optimal utility is positive. Since (162) implies $u_i(\mathbf{0}) = 0$ and $\mathbf{p}^* > \mathbf{0}$, we have $b_i > 0$. Feasibility and $\mathbf{p}^* \in \Delta_n$ imply $\sum_i b_i \le 1$, hence $b_i \le 1$. The bundle $\frac{w_i}{b_i}\mathbf{x}_i^{\mathrm{opt}}$ costs w_i at $\mathbf{p}^*$, so demand optimality implies

$$u_i\left(\frac{w_i}{b_i}\mathbf{x}_i^{\mathrm{opt}}\right) \le u_i(\mathbf{x}_i(\mathbf{p}^*)). \tag{185}$$

By Assumption I.1 and (185),

$$\begin{aligned} \log \frac{u_i(\mathbf{x}_i^{\mathrm{opt}})}{u_i(\mathbf{x}_i(\mathbf{p}^*))} &\le \begin{cases} \overline{r} \log \frac{b_i}{w_i}, & b_i \ge w_i, \\ \underline{r} \log \frac{b_i}{w_i}, & b_i < w_i \end{cases} \\ &= \underline{r} \log \frac{b_i}{w_i} + (\overline{r} - \underline{r}) \max\left\{0, \log \frac{b_i}{w_i}\right\}. \end{aligned} \tag{186}$$

Since $b_i \le 1$ and $w_i \le 1$, it follows that $\max\{0, \log \frac{b_i}{w_i}\} \le \log \frac{1}{w_i}$, and therefore

$$\begin{aligned} \log \frac{\mathrm{NSW}^{\mathrm{opt}}}{\mathrm{NSW}(\mathbf{x}(\mathbf{p}^*))} &= \sum_{i=1}^m w_i \log \frac{u_i(\mathbf{x}_i^{\mathrm{opt}})}{u_i(\mathbf{x}_i(\mathbf{p}^*))} \\ &\le \underline{r} \sum_{i=1}^m w_i \log \frac{b_i}{w_i} + (\overline{r} - \underline{r}) H(\mathbf{w}) \\ &\overset{(a)}{\le} \underline{r} \log \left(\sum_{i=1}^m b_i\right) + (\overline{r} - \underline{r}) H(\mathbf{w}) \\ &\le (\overline{r} - \underline{r}) H(\mathbf{w}). \end{aligned} \tag{187}$$

Here (a) is Jensen's inequality for the concave logarithm with weights $\mathbf{w}$. As in Garg et al. (2026, Theorem 3.1), the inequality $\log t/t \le 1/e$ for $t \ge 1$ implies

$$\sum_{i=1}^{m} w_i \max\left\{0, \log \frac{b_i}{w_i}\right\} \le \frac{1}{e}\sum_{i=1}^{m} b_i \le \frac{1}{e}. \tag{188}$$

By the Jensen step above, (186) also yields

$$\log \frac{\mathrm{NSW}^{\mathrm{opt}}}{\mathrm{NSW}(\mathbf{x}(\mathbf{p}^*))} \le \underline{r} \log\left(\sum_{i=1}^{m} b_i\right) + \frac{\overline{r}-\underline{r}}{e} \le \frac{\overline{r}-\underline{r}}{e}. \tag{189}$$

The definition of M^* and (175) imply

$$\widehat{\mathbf{x}}_i(\mathbf{p}^*) \ge \frac{1}{M^*}\mathbf{x}_i(\mathbf{p}^*), \qquad i \in [m]. \tag{190}$$

Monotonicity, Assumption I.1, and $M^* \ge 1$ imply

$$\begin{aligned} u_i(\widehat{\mathbf{x}}_i(\mathbf{p}^*)) &\ge u_i\left(\frac{1}{M^*}\mathbf{x}_i(\mathbf{p}^*)\right) \\ &\ge (M^*)^{-\overline{r}} u_i(\mathbf{x}_i(\mathbf{p}^*)). \end{aligned} \tag{191}$$

Combining (187), (189), and (191), with $\sum_i w_i = 1$, yields

$$\begin{aligned} \Delta_{\log} &= \log \frac{\mathrm{NSW}^{\mathrm{opt}}}{\mathrm{NSW}(\mathbf{x}(\mathbf{p}^*))} + \sum_{i=1}^{m} w_i \log \frac{u_i(\mathbf{x}_i(\mathbf{p}^*))}{u_i(\widehat{\mathbf{x}}_i(\mathbf{p}^*))} \\ &\le (\overline{r}-\underline{r}) \min\{H(\mathbf{w}), 1/e\} + \overline{r} \log M^* \\ &\le (\overline{r}-\underline{r}) \min\{H(\mathbf{w}), 1/e\} + \overline{r} \log\left(1 + \frac{\eta^* + \xi^*}{p^*_{\min}}\right). \end{aligned} \tag{192}$$

The first upper bound in (192), followed by exponentiation, proves the upper bounds in (179). On the event in (174), we have $\eta^* \le \eta$, $\xi^* \le \epsilon$, and $p^*_{\min} \ge \underline{p}$, proving the upper bounds in (180). Exponentiating this bound yields

$$\frac{\mathrm{NSW}(\widehat{\mathbf{x}}(\mathbf{p}^*))}{\mathrm{NSW}^{\mathrm{opt}}} \ge e^{-(\overline{r}-\underline{r})\min\{H(\mathbf{w}),1/e\}} \left(\frac{\underline{p}}{\underline{p}+\eta+\epsilon}\right)^{\overline{r}}. \tag{193}$$

Equation (160) then implies the upper bound in (181). Feasibility of (175) implies $\Delta_{\log} \ge 0$ and $\Delta_{\mathrm{NSW}} \ge 0$. ■

Remark 8. *The demand reports at deployment determine M^*, so the realized bound* (179) *holds without Assumptions I.2 and I.3; only the sample-based bound* (181) *uses them. For degree-one homogeneous utilities, $\underline{r} = \overline{r} = 1$, so Theorem I.1 recovers*

$$\Delta_{\log} \le \log\left(1 + \frac{\eta+\epsilon}{\underline{p}}\right), \qquad \Delta_{\mathrm{NSW}} \le \frac{\eta+\epsilon}{\underline{p}+\eta+\epsilon}. \tag{194}$$

For GES utilities, (164) *makes the guarantee applicable to the non-homothetic experiments in Section 5.2. For non-homogeneous utilities, an exact equilibrium need not be proportionally fair, even when prediction and equilibrium errors vanish. The first term bounds this discrepancy independently of the number of agents.*