

Robust Market Equilibria with Uncertain Preferences

Riley Murray,^{1,3} Christian Kroer,^{1,4} Alex Peysakhovich,² Parikshit Shah¹

¹Facebook Core Data Science, ²Facebook Artificial Intelligence Research

³California Institute of Technology, ⁴Columbia University

rmurray@caltech.edu, christian.kroer@columbia.edu, alex.peys@gmail.com, parikshit@fb.com

Abstract

The problem of allocating scarce items to individuals is an important practical question in market design. An increasingly popular set of mechanisms for this task uses the concept of market equilibrium: individuals report their preferences, have a budget of real or fake currency, and a set of prices for items and allocations is computed that sets demand equal to supply. An important real world issue with such mechanisms is that individual valuations are often only imperfectly known. In this paper, we show how concepts from classical market equilibrium can be extended to reflect such uncertainty. We show that in linear, divisible Fisher markets a robust market equilibrium (RME) always exists; this also holds in settings where buyers may retain unspent money. We provide theoretical analysis of the allocative properties of RME in terms of envy and regret. Though RME are hard to compute for general uncertainty sets, we consider some natural and tractable uncertainty sets which lead to well behaved formulations of the problem that can be solved via modern convex programming methods. Finally, we show that very mild uncertainty about valuations can cause RME allocations to outperform those which take estimates as having no underlying uncertainty.

1 Introduction

A key problem in market design is ‘who gets what’ (Roth 2015). An important mechanism for the allocation of scarce items to multiple individuals is the use of market equilibrium. In these mechanisms individuals have preferences over scarce items and budgets of money. Prices are set for items such that demand of individuals equals the supply of items. A major issue in practice with such mechanisms is that the utility of an individual is often not known exactly (sometimes even to the individuals themselves). In this paper, we take up the question of computing a robust market equilibrium which takes this imperfect information into account.

Market equilibrium-based allocations are increasingly prevalent in real world mechanisms and robustness is an important issue in many of these applications.

In the allocation of courses to students at business schools students report preferences over courses, are given a budget of ‘fake’ currency, and are allocated the courses they receive

in the market equilibrium which matches supply of courses to demand (Budish and Cantillon 2012). Here students may not be fully aware of their own exact valuations for courses.

In online advertising advertisers report valuations for various impressions, enter budgets, and an auction mechanism again sets supply equal to demand - finding an equilibrium of the market (Balseiro, Besbes, and Weintraub 2015; Balseiro et al. 2017; Balseiro and Gur 2017). In practice advertisers’ bids in the auction (which are equivalent to a paced version of their valuation in equilibrium (Borgs et al. 2007; Conitzer et al. 2018; 2019)) come from a combination of a per-interaction valuation and a machine learning model which predicts the probability that a given user will interact with that ad. For example, an advertiser might say they are willing to pay \$1 per click and a model could predict that a user will click a particular ad with probability .1 – this would lead to a bid of 10 cents on the ad when combined with the valuations. Here, again, the valuation is imperfectly known as both the advertiser’s willingness to pay and, more so, the click prediction model, come with noise.

We begin with the simplest workhorse model in this literature: linear Fisher markets. In such markets there are buyers and items. Buyers have finite budgets. Items are scarce and divisible. Buyers have utility for each item and value a bundle as the utility-weighted sum of the items in the bundle. We consider both standard Fisher markets where buyer budgets do not have any use outside the market, and quasi-Fisher markets where leftover money has a fixed value for each buyer. Equilibria in such markets include allocations to individuals and prices for items, so that buyer demands sum up to supply and allocations satisfy buyer demands.

In this article, we show how to extend market equilibrium to the notion of robust market equilibrium (RME). Here buyers do not have point-valued utilities for each item, but rather have ‘uncertainty sets’ of possible valuations. In this case a buyer wants to purchase bundles that optimize their worst case utility given prices. This can reflect risk aversion on the part of the buyers (buyers are actually not perfectly aware of item utility) or ignorance on the part of the market designer if the market equilibrium is simply used as an allocative mechanism (e.g. CEEI). RME can be partly be viewed as a model of classical equilibrium with a particular class of nonlinear utility functions. Indeed, the *existence* of an RME is no different than the existence of a classical equilibrium

with appropriate utilities; the main point of departure is how the resulting equilibria are measured. We extend classical notions of envy and regret to reflect uncertainty in buyer valuations, and we give bounds on these quantities as functions of the size of the uncertainty sets.

It is well-known that computing market equilibria can be computationally intractable for certain market models. One of the earliest positive results in this area is the celebrated Eisenberg-Gale convex program for Fisher markets (Eisenberg and Gale 1959); this convex program was first proposed for linear utility functions, but was later shown to work for any utilities which are homogeneous of degree one (Eisenberg 1961). In the past 20 years, this convex programming approach has been extended to handle more general Fisher markets such as spending constraint utilities (Birnbbaum, Devanur, and Xiao 2011), utility-capped buyers, quasi-linear buyers (Chen, Ye, and Zhang 2007; Cole et al. 2017), and Fisher markets with transaction costs (Chakraborty, Devanur, and Karande 2010).

Over the course of this article, we show how techniques from the robust optimization literature can be combined with the Eisenberg-Gale convex program to compute RME in both Fisher and quasi-Fisher markets. In principle, this means the existence of polynomial-time algorithms for computing Fisher or quasi-Fisher RME reduces to the existence of a polynomial-time algorithm for evaluating robust utilities. In practice, we need to represent the value of a robust utility as a tractable conic optimization problem. We propose two types of uncertainty sets that lead to convex programs which can be efficiently solved using modern solvers, and we apply these proposals to study RME in real datasets.

To the best of our knowledge, we are the first to study robust market equilibria. However, the topic of robust variants of optimization problems has been studied extensively (Bent-Tal and Nemirovski 2002; Bertsimas and Sim 2004; Bent-Tal, El Ghaoui, and Nemirovski 2009; Bertsimas, Brown, and Caramanis 2011). There, some *nominal* mathematical program is given, and then the robust variant of the program requires that the constraints hold for every instantiation of parameters from some uncertainty set. Our robust market equilibrium approach can be viewed as the natural uncertainty parameterization of Eisenberg-Gale style convex programs. A similar uncertainty-parametrization of utilities was considered by Aghassi and Bertsimas (2006) in the context of game-theoretic equilibrium where they provide a robust analogue of Bayesian equilibrium. Robust game-theoretic equilibria have also been considered in the context of counterfactual prediction (Peysakhovich, Kroer, and Lerer 2019). Finally, there is literature on *robust mechanism design* (Bergemann and Morris 2005; Lopomo, Rigotti, and Shannon 2018; Albert et al. 2017), where the goal is to design mechanisms that are robust either to uncertainty about the distribution over agent payoffs, or the belief that an agent holds about the types of other agents. Due to the close relationship between mechanism design and market equilibria it would be interesting to understand what kind of robustness properties RME has with respect to traditional robust mechanism design objectives.

2 Fisher Markets

Consider a market where n buyers compete for m divisible goods, each of unit supply. Each buyer brings a budget b_i of currency to this market, and a utility function u_i over allocations of the m goods. The problem of market equilibria is to determine prices for goods and allocations of goods to buyers which satisfy both supply constraints, and certain optimality conditions for each buyer. The supply constraints do not depend on further details of the market. If X is a nonnegative matrix where x_{ij} is the allocation of good j to buyer i , the allocation is **feasible** if $\sum_{i=1}^n x_{ij} \leq 1$ for all goods j .

The precise form of the buyer optimality conditions depends on whether or not money has value outside the market. Suppose that the market has assigned each good j a price p_j , and assemble these prices into a vector \mathbf{p} in \mathbb{R}_+^m . When money has no intrinsic value, we are in a **Fisher market**. A **market equilibrium** is a feasible allocation X and a set of prices \mathbf{p} such that each buyer prefers what they are allocated over anything else they could afford:

$$\mathbf{x}_i \in \operatorname{argmax}\{u_i(\mathbf{z}) \mid \mathbf{z} \in \mathbb{R}_+^m \text{ and } \mathbf{p}^* \cdot \mathbf{z} \leq b_i\}. \quad (1)$$

When money has intrinsic value, buyer utilities become $(\mathbf{z}, s) \mapsto u_i(\mathbf{z}) + s$, where s is the buyer's retained budget. Here the buyer optimality conditions for market equilibria are that

$$(\mathbf{x}_i, r_i) \in \operatorname{argmax}\{u_i(\mathbf{z}) + s \mid \mathbf{z} \in \mathbb{R}_+^m, s \in \mathbb{R}_+, \text{ and } \mathbf{p} \cdot \mathbf{z} + s \leq b_i\} \quad (2)$$

for the specific value $r_i = b_i - \mathbf{p} \cdot \mathbf{x}_i$. Throughout this article we call this model a **quasi-Fisher market**, because the effective utility functions $(\mathbf{z}, s) \mapsto u_i(\mathbf{z}) + s$ are broadly known as *quasilinear utilities* in the economics literature.

Linear utility functions offer the simplest model for buyer preferences in market equilibrium problems. In a linear model, each buyer possesses a valuation vector $\mathbf{v}_i \in \mathbb{R}_+^m$, and assigns utility $u_i(\mathbf{z}) = \mathbf{v}_i \cdot \mathbf{z}$ to a bundle of goods \mathbf{z} . There is one serious drawback to using linear utilities: in large markets, it is not realistic to assume that all \mathbf{v}_i 's are known exactly—either by buyers, or by a market maker.

We thus propose a model where each buyer has an associated **uncertainty set of valuations** $\mathcal{V}_i \subset \mathbb{R}_+^m$; an uncertainty set gives rise to a **robust utility**

$$u_i(\mathbf{z}) = \min\{\mathbf{z} \cdot \mathbf{v} \mid \mathbf{v} \in \mathcal{V}_i\}. \quad (3)$$

It is easy to see that for all nonempty \mathcal{V}_i , these robust utilities are concave, and positively homogeneous of degree 1. We assume these uncertainty sets are nonempty, compact, and satisfy $u_i(\mathbf{1}) > 0$. With regards to \mathcal{V}_i being nonempty, we will often find it useful to suppose that each buyer has a distinguished **nominal valuation** $\hat{\mathbf{v}}_i \in \mathcal{V}_i$. We often call the function $\mathbf{z} \mapsto \hat{\mathbf{v}}_i \cdot \mathbf{z}$ the **nominal utility** of buyer i .

2.1 Computing market equilibria

In a foundational result, Eisenberg and Gale (1959) showed that equilibria for Fisher markets with linear utility functions can be computed by solving a particular centralized convex program. In a follow-up work Eisenberg (1961) showed the

same convex program can be used to compute Fisher-market equilibria whenever utilities are concave, and homogeneous of degree one. Much later, Chen, Ye, and Zhang (2007) gave a simpler proof for Eisenberg’s 1961 result, which extended naturally to compute equilibria in quasi-Fisher markets.¹

Chen et al.’s result for quasi-Fisher markets requires a very minor modification to the convex program originally proposed by Eisenberg and Gale. We can state both convex programs by considering a parametric optimization problem: let $Q \in \{0, 1\}$ be a parameter, where $Q = 1$ indicates a quasi-Fisher market model, and $Q = 0$ indicates a Fisher market model. The **Eisenberg-Gale convex program** is

$$\begin{aligned} \max \quad & \sum_{i=1}^n b_i \log(t_i) - Q r_i & (4) \\ \text{s.t.} \quad & (\mathbf{x}_i, t_i, r_i) \in \mathbb{R}_+^{m+2} & \forall i \in [n], \\ & t_i \leq u_i(\mathbf{x}_i) + Q r_i & \forall i \in [n], \\ & \sum_{i=1}^n x_{ij} \leq 1 & \forall j \in [m]. \end{aligned}$$

The results of Eisenberg and Chen et al. ensure that as long as u_i are concave and homogeneous of degree one, then for $\mathbf{p} \in \mathbb{R}_+^m$ as the vector of dual variables to the capacity constraints, the matrix \mathbf{X} whose rows are \mathbf{x}_i forms an equilibrium allocation with respect to prices \mathbf{p} .

Because we defined our robust utilities as the minimum over a set of linear functions, we get that robust utilities are homogeneous with degree 1 and concave. This establishes the following proposition:

Proposition 1. *A solution to (4) with $Q = 0$ produces an equilibrium for the Fisher market where u_i are given by (3). Optimal prices \mathbf{p} and allocations \mathbf{X} satisfy the following properties:*

1. Each \mathbf{x}_i is in the demand set of buyer i , i.e. (1) holds.
2. Every item j with $p_j > 0$ clears the market: $\sum_i x_{ij} = 1$.
3. For every buyer i , there exists a $\mathbf{v}_i^* \in \operatorname{argmin}_{\mathbf{v} \in \mathcal{V}_i} \mathbf{v} \cdot \mathbf{x}_i$ for which allocated items have the same bang per buck:

$$\text{if } x_{ij}, x_{ik} > 0 \text{ then } v_{ij}^*/p_j = v_{ik}^*/p_k.$$

The first property in Proposition 1 is about individual optimality; every buyer seeking to optimize her own utility in the robust sense will find the allocations in her demand set. The second property shows market clearing. The third property shows that the worst case utilities are in fact attained in the allocation, and that an “equal rates” condition holds at one of these worst-case utility vectors.

Proposition 2. *A solution to (4) with $Q = 1$ produces an equilibrium for the quasi-Fisher market with u_i given by (3). Analogous conditions to those in Proposition 1 hold for prices \mathbf{p} and allocations \mathbf{X} . Furthermore, if β denotes optimal dual variables to the utility hypograph constraints, then for all buyers i*

1. $\beta_i \leq 1$, and $\beta_i = 1$ whenever $\mathbf{p} \cdot \mathbf{x}_i < b_i$.
2. There is a $\mathbf{v}_i^* \in \operatorname{argmin}_{\mathbf{v} \in \mathcal{V}_i} \mathbf{v} \cdot \mathbf{x}_i$ so that $\beta_i \leq p_j/v_{ij}^*$ for all j , and $\beta_i = p_j/v_{ij}^*$ whenever $x_{ij} > 0$.

¹In situ, Chen et al. refer to quasi-Fisher markets as *mixed Fisher Arrow-Debreu markets*.

The additional conditions concern interpretation of the dual variables β_i as *pacing multipliers* in a first-price auction (Conitzer et al. 2019). The allocation mechanism may be viewed as a first price auction where buyer i competes for item j with a modified or “paced” bid of $\beta_i v_{ij}$. The pacing multipliers are no greater than one (ensuring that no buyer buys a good when its price exceeds its value), there is no unnecessary pacing, and the robust bang-per-buck is equal for all allocated items to buyer i .

Propositions 1 and 2 argue that a solution to (4) leads to a reasonable solution concept: for risk averse agents that are seeking to robustly maximize their utilities in the face of the market uncertainty, it produces allocations in their individual demand sets. Moreover, it also leads to market clearing, which is desirable for the market designer. Finally, the solution is intuitive, and in fact corresponds to a “standard” market equilibrium with respect to a set of attained (worst case) realizations of the utilities.

From a theoretical perspective, computing market equilibria with robust utilities is no harder than solving an appropriate convex program. Unfortunately, not all convex programs are tractable; the potential stumbling block in our case is the need to represent the robust utilities via hypographs

$$H_i \doteq \{(t, \mathbf{z}) \mid t \leq u_i(\mathbf{z})\}$$

using a limited library of convex constraints on t and \mathbf{z} . When the uncertainty sets \mathcal{V}_i have an explicit convex description, the H_i can be represented by appealing to convex duality. For example, if $\mathcal{V}_i = \{\mathbf{v} \mid \|\mathbf{v} - \hat{\mathbf{v}}_i\| \leq \epsilon\}$ for the nominal valuation $\hat{\mathbf{v}}_i$ and some reference norm $\|\cdot\|$, the hypograph can be represented as

$$H_i = \{(t, \mathbf{z}) \mid t + \epsilon \|\mathbf{z}\|^* \leq \hat{\mathbf{v}}_i \cdot \mathbf{z}\}$$

where $\|\cdot\|^*$ is the dual norm to $\|\cdot\|$. In cases such as this, we can appeal to standard convex programming solvers, which allow only specific convex operations.

In principle, the \mathcal{V}_i need not be convex. This is for the following reason: if “conv” is the operator that computes a set’s convex hull, then

$$\min\{\mathbf{z} \cdot \mathbf{v} \mid \mathbf{v} \in \mathcal{V}_i\} = \min\{\mathbf{z} \cdot \mathbf{v} \mid \mathbf{v} \in \operatorname{conv} \mathcal{V}_i\}$$

for all \mathbf{z} in \mathbb{R}^m . Thus we may replace nonconvex uncertainty sets by their convex hulls without loss of generality. The challenge of nonconvex uncertainty sets then reduces to the problem of efficiently representing their convex hulls.

2.2 A dual formulation

The literature on theoretical analysis of market equilibrium often benefits from analyzing primal and dual formulations in conjunction with one another. Given the nature of our utility functions, it is reasonably straightforward to compute the dual by appealing to Fenchel duality. The drawback to working with a Fenchel dual is that it can become harder to interpret dual variables. To assist others in future theoretical analysis of market equilibria with robust utility functions, we provide the following result (the proof of which can be found in a separate supplemental document, along with all other omitted proofs of claims in this article):

Proposition 3. Let u_i be given by (3) where $\mathcal{V}_i \subset \mathbb{R}_+^m$ are nonempty compact convex sets, and $u_i(\mathbf{1}) > 0$. It can be shown that the following problem is a dual to (4)

$$\begin{aligned} \min \quad & \mathbf{p} \cdot \mathbf{1} - \sum_{i=1}^n \{b_i + b_i \log(\beta_i/b_i)\} \\ \text{s.t.} \quad & \mathbf{p} \in \mathbb{R}_+^m, \beta \in \mathbb{R}_+^n \\ & \mathbf{Q}[1 - \beta] \geq \mathbf{0} \\ & \mathbf{p} \geq \beta_i \mathbf{v}_i \quad \forall i \in [n] \\ & \mathbf{v}_i \in \mathcal{V}_i \quad \forall i \in [n]. \end{aligned}$$

The formulation is nonconvex as stated, due to bilinear inequality constraints $\mathbf{p} \geq \beta_i \mathbf{v}_i$ in the variables β_i and \mathbf{v}_i . This bilinearity can be represented in a convex way by using perspective functions. For each $i \in [n]$, set $\mathcal{W}_i = \mathbb{R}_+^m + \mathcal{V}_i$.² The \mathcal{W}_i are convex sets, and so their indicator functions

$$\delta_{\mathcal{W}_i}(\mathbf{y}) = \begin{cases} 0 & \text{if } \mathbf{y} \in \mathcal{W}_i \\ +\infty & \text{otherwise} \end{cases}$$

are also convex. The bilinear constraints can thus be written

$$\beta_i \delta_{\mathcal{W}_i}(\mathbf{p}/\beta_i) \leq 0 \quad \forall i \in [n].$$

These constraints are convex, since the mappings on the left-hand-side of the inequalities are the perspectives of the convex functions $\delta_{\mathcal{W}_i}$. It can be shown that this formulation (in terms of perspectives-of-indicators) is in fact precisely what results when applying Fenchel duality to (4) directly.

3 Properties of Robust Equilibria

On the surface, the market equilibria considered in this article are no different than classical market equilibria with a particular choice of nonlinear utility function. Thus there are standard properties of market equilibria which are naturally satisfied in our setting. For example, if (\mathbf{X}, \mathbf{p}) form an equilibrium in a Fisher market, and all budgets $b_i = 1$, then the allocations are *envy-free* in the sense that for all buyers i

$$u_i(\mathbf{x}_i) \geq u_i(\mathbf{x}_{i'}) \text{ for all } i' \text{ in } [n].$$

Similarly, the allocations have *no regret*, in that for all i

$$R_i(\mathbf{y}, \mathbf{p}) \doteq \max\{u_i(\mathbf{z}) - u_i(\mathbf{y}) \mid \mathbf{z} \in \mathbb{R}_+^m, \mathbf{p} \cdot \mathbf{z} \leq b_i\}$$

satisfies $R_i(\mathbf{x}_i, \mathbf{p}) = 0$. These classical measures of fairness are important, but they do not reflect our reason for choosing the robust utilities $u_i(\mathbf{z}) = \min\{\mathbf{z} \cdot \mathbf{v} \mid \mathbf{v} \in \mathcal{V}_i\}$ in the first place.

We assume that once a buyer has received a bundle of goods \mathbf{x}_i , some $\mathbf{v}_i \in \mathcal{V}_i$ instantiates as the buyer's true valuation. This can occur, for example, if uncertainty came from the fact that buyers did not have perfect clarity into what they valued during evaluation time but gain this information when they actually receive the good. In this model it is natural to evaluate metrics such as envy, regret, or Nash Social Welfare with respect to these **realizations** of buyer uncertainty. We focus on the scenarios when the valuation realizes to its nominal value, and when it realizes adversarially.

²Addition here denotes the Minkowski sum.

3.1 Adversarially-robust envy

Let \mathbf{X} denote a matrix whose rows \mathbf{x}_i are equilibrium allocations with respect to prices \mathbf{p} and robust utilities u_i .

Also let $r_i = b_i - \mathbf{p} \cdot \mathbf{x}_i$ denote the retained budget for each buyer $i \in [n]$ (understanding that $r_i = 0$ in Fisher markets). In these terms, robust envy is

$$\begin{aligned} \mathcal{E}_i(\mathbf{X}, \mathbf{p}) &= \max(\mathbf{v} \cdot \mathbf{x}_{i'} + r_{i'}) - (\mathbf{v} \cdot \mathbf{x}_i + r_i) \\ \text{s.t.} \quad & \mathbf{v} \in \mathcal{V}_i, i' \in [n]. \end{aligned}$$

In order to analyze robust envy, it is helpful to frame things in terms of the underlying robust utility functions. One may verify that the following identity holds

$$\mathcal{E}_i(\mathbf{X}, \mathbf{p}) = \max_{i' \in [n]} \{r_{i'} - r_i - u_i(\mathbf{x}_i - \mathbf{x}_{i'})\}.$$

It is therefore evident that we can efficiently compute robust envy whenever we can compute the utility functions u_i .

3.2 Adversarially-robust regret

Given prices \mathbf{p} , the i^{th} buyer takes an action by drawing from the following polytope

$$Z_i(\mathbf{p}) = \{\mathbf{z} \mid \mathbf{0} \leq \mathbf{z} \leq \mathbf{1}, \mathbf{p} \cdot \mathbf{z} \leq b_i\}.$$

Then given an allocation \mathbf{x}_i , we thus define a buyer's robust regret as

$$\begin{aligned} \mathcal{R}_i(\mathbf{x}_i, \mathbf{p}) &= \max[\mathbf{v} - \mathbf{Q}\mathbf{p}] \cdot \mathbf{z} - [\mathbf{v} - \mathbf{Q}\mathbf{p}] \cdot \mathbf{x}_i \\ \text{s.t.} \quad & \mathbf{v} \in \mathcal{V}_i, \mathbf{z} \in Z_i(\mathbf{p}) \end{aligned}$$

– where “Q” is 1 if the buyer is in a quasi-Fisher market, and 0 if otherwise. Mirroring our approach to robust-envy, robust-regret can be expressed with the robust utility functions. Specifically,

$$\begin{aligned} \mathcal{R}_i(\mathbf{x}_i, \mathbf{p}) &= \max \mathbf{Q}\mathbf{p} \cdot [\mathbf{x}_i - \mathbf{z}] - u_i(\mathbf{x}_i - \mathbf{z}) \\ \text{s.t.} \quad & \mathbf{z} \in Z_i(\mathbf{p}). \end{aligned}$$

Robust regret is harder to compute than robust envy. The issue is that the objective appearing in the second expression for $\mathcal{R}_i(\mathbf{x}_i, \mathbf{p})$ is convex, rather than concave. From a complexity perspective, it is very difficult to maximize convex functions. Our one reason for hope is that we are trying to maximize a convex function over the *polytope* $Z_i(\mathbf{p})$. In this specific case we may use the fact that the maximum of a real-valued convex function over a polytope is always attained at one of the polytope's extreme points. Thus in principle, it is possible to compute $\mathcal{R}_i(\mathbf{x}_i, \mathbf{p})$ by evaluating the objective function at each extreme point of $Z_i(\mathbf{p})$, and taking the largest value. This approach may be viable in Fisher markets for buyers that are heavily budget constrained, since if $b_i \leq p_j$ for all goods j , then the extreme points of $Z_i(\mathbf{p})$ are $\mathbf{z}^{(0)} = \mathbf{0}$, and $\mathbf{z}^{(j)} = b_i/p_j \mathbf{e}_j$ for all $j \in [m]$.³ In other contexts the extreme-point approach will likely be impractical. For example, in Fisher markets where prices are low compared to a buyer's budget, the set $Z_i(\mathbf{p})$ can contain as many as 2^m extreme points.

³Here, \mathbf{e}_j denotes the j^{th} standard basis vector in \mathbb{R}^m .

3.3 Bounds in Fisher markets

In this section we describe elementary ways to bound robust envy and robust regret for Fisher markets.

Lemma 1. *Suppose the ℓ_1 -width of \mathcal{V}_i is bounded above by w . Then for all \mathbf{v} in \mathcal{V}_i and all \mathbf{z} with $\|\mathbf{z}\|_\infty \leq 1$, we have*

$$-w + \mathbf{v} \cdot \mathbf{z} \leq u_i(\mathbf{z}) \leq \mathbf{v} \cdot \mathbf{z} + w.$$

Proof. Fix \mathbf{v} in \mathcal{V}_i and \mathbf{z} with $\|\mathbf{z}\|_\infty \leq 1$. Since \mathcal{V}_i is compact, there exists a $\mathbf{v}^* \in \mathcal{V}_i$ so that $u_i(\mathbf{z}) = \mathbf{v}^* \cdot \mathbf{z}$. Since the ℓ_1 -width of \mathcal{V}_i is bounded by w , we have that $\Delta \mathbf{v} = \mathbf{v}^* - \mathbf{v}$ satisfies $\|\Delta \mathbf{v}\|_1 \leq w$. Writing $u_i(\mathbf{z}) = \mathbf{v} \cdot \mathbf{z} + \Delta \mathbf{v} \cdot \mathbf{z}$, we can invoke duality between the ℓ_1 and ℓ_∞ norms to see that $|\Delta \mathbf{v} \cdot \mathbf{z}| \leq w$. The result follows. \square

Lemma 1 is useful because the natural domain for u_i in robust regret and envy is the unit ℓ_∞ ball.

Proposition 4. *Suppose the ℓ_1 -width of \mathcal{V}_i is bounded above by w . Then robust regret is bounded by $2w$. If all budgets are equal then robust envy is bounded by $2w$.*

Proof. Let \mathbf{X}, \mathbf{p} be a robust market equilibrium. Let i^*, \mathbf{v}^* be the maximizers of robust envy for buyer i . We then have

$$\begin{aligned} \mathcal{E}_i(\mathbf{X}, \mathbf{p}) &= (\mathbf{v}^* \cdot \mathbf{x}_{i^*} + r_{i^*}) - (\mathbf{v}^* \cdot \mathbf{x}_i + r_i) \\ &\leq (w + u_i(x_{i^*}) + r_{i^*}) - (-w + u_i(x_i) + r_i) \\ &= 2w + (u_i(x_{i^*}) + r_{i^*}) - (u_i(x_i) + r_i) \\ &\leq 2w, \end{aligned}$$

where the first inequality follows by Lemma 1 (since supplies are 1), and the second inequality follows by the no-envy property since \mathbf{X}, \mathbf{p} constitute a market equilibrium with respect to the robust utilities and budgets are equal.

The proof for robust regret is analogous, but using the fact that $u_i(\mathbf{x}_i) - \mathbf{Q} \mathbf{p} \cdot \mathbf{x}_i \geq u_i(\mathbf{z}) - \mathbf{Q} \mathbf{p} \cdot \mathbf{z}$ where \mathbf{z} is the allocation for robust regret. \square

While the above bounds hold uniformly for every agent i , in the specific case when the ℓ_∞ -width of the uncertainty sets are small one can also obtain a bound on the *average* robust regret of all the buyers.

Proposition 5. *Suppose that for each buyer $i \in [n]$, we have $|v_{ij} - v'_{ij}| \leq R$ for all $j \in [m]$ and all $\mathbf{v}_i, \mathbf{v}'_i \in \mathcal{V}_i$. Then the following holds:*

$$\frac{1}{n} \sum_{i=1}^n \text{Robust-Regret}_i \leq \frac{2Rm}{n}.$$

In a scarce market where the number of buyers far exceeds the number of items, the bound in Proposition 5 guarantees that the average regret is small.

4 Concrete Models for Buyer Uncertainty

We consider the case of low rank markets (Kroer et al. 2019; Peysakhovich and Kroer 2019). In these markets, the valuations individuals place on items are not independent, and can be predicted from one another; this is common in most real-world allocation problems (e.g. valuations for courses are correlated within person, ratings of movies are correlated

within person, willingness to pay for different impressions are correlated within advertiser).

Let $\widehat{\mathbf{V}}$ denote our estimated valuation matrix for the Fisher market, i.e. \hat{v}_{ij} is the value that individual i places on item j . The low rank model assumes that $\widehat{\mathbf{V}} = \widehat{\Theta} \widehat{\Phi}^\top$ where $\widehat{\Theta}$ and $\widehat{\Phi}$ are embedding matrices for individuals and items respectively, of sizes $n \times d$ and $m \times d$. The problem is typically interesting in the regime where $d \ll n, m$.

In such a model, there are a few natural sources of uncertainty. For example, in practice, individual and item embeddings would typically be estimated using some matrix completion procedure. The matrix completion procedure may introduce errors (relative to the ground-truth completed matrix), and these errors induce uncertainty on factors $\widehat{\Theta}$ or $\widehat{\Phi}$. There is also the possibility that the low-rank assumption does not well reflect reality; this would correspond to uncertainty in small-norm but high-rank perturbations to $\widehat{\mathbf{V}}$.

When the embedding model is imperfect and subject to uncertainty, a natural description one may consider is the **joint uncertainty model** defined by the set $\mathcal{V}_i^J(\epsilon_1, \epsilon_2)$ as:

$$\begin{aligned} \{\boldsymbol{\theta} \widehat{\Phi}^\top \mid \boldsymbol{\theta} \in \mathbb{R}^{1 \times d}, \boldsymbol{\theta} \widehat{\Phi}^\top \geq \mathbf{0}, \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_i\|_p \leq \epsilon_1 \|\hat{\boldsymbol{\theta}}_i\|_p, \\ \|\Phi - \widehat{\Phi}\|_q \leq \epsilon_2 \|\widehat{\Phi}\|_q, \boldsymbol{\theta} \widehat{\Phi}^\top \cdot \mathbf{1} = \hat{\boldsymbol{\theta}}_i \widehat{\Phi}^\top \cdot \mathbf{1}\}. \end{aligned}$$

This model captures uncertainty in both the buyer-side and item-side embedding representations. The model parameters include a pair of norms (the p -norm for the vector quantity, and the q -norm for the matrix quantity),⁴ as well as relative radii $\epsilon_1, \epsilon_2 \in (0, 1)$.

Certain special cases of the joint uncertainty model are interesting to consider. The **direct model** for uncertainty under such parameters is realized when $\epsilon_2 = 0$, and $\widehat{\Phi} = I$, the identity matrix. In this situation, we have:

$$\begin{aligned} \mathcal{V}_i^d(p, \epsilon) &= \{\mathbf{v} \mid \mathbf{v} \in \mathbb{R}^m, \mathbf{v} \geq \mathbf{0}, \\ &\|\mathbf{v} - \hat{\mathbf{v}}_i\|_p \leq \epsilon \|\hat{\mathbf{v}}_i\|_p \\ &\mathbf{v} \cdot \mathbf{1} = \hat{\mathbf{v}}_i \cdot \mathbf{1}\}. \end{aligned}$$

When $\epsilon_2 = 0$ but $\widehat{\Phi}$ is arbitrary, we obtain the **buyerside model**, which is described by:

$$\begin{aligned} \mathcal{V}_i^b(p, \epsilon) &= \{\boldsymbol{\theta} \widehat{\Phi}^\top \mid \boldsymbol{\theta} \in \mathbb{R}^{1 \times d}, \boldsymbol{\theta} \widehat{\Phi}^\top \geq \mathbf{0}, \\ &\|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_i\|_p \leq \epsilon \|\hat{\boldsymbol{\theta}}_i\|_p, \\ &\boldsymbol{\theta} \widehat{\Phi}^\top \cdot \mathbf{1} = \hat{\boldsymbol{\theta}}_i \widehat{\Phi}^\top \cdot \mathbf{1}\}. \end{aligned}$$

The linear equality constraints in these definitions require that we conserve utility about the vector of all ones (“ $\mathbf{1}$ ”); there are several compelling reasons for doing this. First, it enforces a kind of regularity condition: assuming $\hat{\mathbf{v}}_i$ is nonzero, the robust utilities will be positive under the uniform allocation $\mathbf{x}_i = \mathbf{1}/m$, regardless of $\epsilon \geq 0$ and $p \geq 1$. Positivity of robust utilities under the uniform allocation ensures that the convex program (4) will always be feasible. The second reason for conserving utility about $\mathbf{1}$ is a concentration argument: although it is unrealistic to presume

⁴The “ p ” in “ p -norm” should not be confused with vectors \mathbf{p} of market prices for goods.

that *all* entries \hat{v}_{ij} are known to high precision, mild statistical assumptions allow us to reliably estimate the sum $\sum_{j=1}^m \hat{v}_{ij} = \hat{\mathbf{v}}_i \cdot \mathbf{1}$. Finally, because all \mathbf{v}_i in the uncertainty sets are elementwise nonnegative, the conservation constraint provides a scale-invariance: $\|\mathbf{v}_i\|_1 = \|\hat{\mathbf{v}}_i\|_1$. This prevents us from thinking that some \mathbf{v}_i in \mathcal{V}_i are more likely than others, simply because they have smaller norm.

We now turn to a deriving a tractable representation for the hypographs of u_i under the direct model. Fix $i \in [n]$, $p \geq 1$, and $\epsilon \in (0, 1)$. For $\boldsymbol{\lambda}$ in \mathbb{R}_+^m and $\boldsymbol{\mu}$ in \mathbb{R} , consider the parameterized Lagrangian

$$\mathcal{L}(\mathbf{v}, \boldsymbol{\lambda}, \boldsymbol{\mu}; \mathbf{z}) = \mathbf{v} \cdot \mathbf{z} + \boldsymbol{\mu}(\mathbf{v} \cdot \mathbf{1} - \hat{\mathbf{v}}_i \cdot \mathbf{1}) - \mathbf{v} \cdot \boldsymbol{\lambda}$$

where \mathbf{v} is restricted to $D_i \doteq \{\mathbf{v} \mid \|\mathbf{v} - \hat{\mathbf{v}}_i\|_p \leq \epsilon \|\hat{\mathbf{v}}_i\|_p\}$. Note how the Lagrangian is convex in \mathbf{v} , concave in $\boldsymbol{\lambda}, \boldsymbol{\mu}$, and how D_i is a compact convex set. We can thus invoke strong duality in a minimax representation for $u_i(\mathbf{z})$:

$$\begin{aligned} u_i(\mathbf{z}) &= \min_{\mathbf{v} \in D_i} \left\{ \max_{\substack{\boldsymbol{\lambda} \geq \mathbf{0} \\ \boldsymbol{\mu} \in \mathbb{R}}} \{\mathcal{L}(\mathbf{v}, \boldsymbol{\lambda}, \boldsymbol{\mu}; \mathbf{z})\} \right\} \\ &= \max_{\substack{\boldsymbol{\lambda} \geq \mathbf{0} \\ \boldsymbol{\mu} \in \mathbb{R}}} \left\{ \min_{\|\Delta \mathbf{v}\|_p \leq \epsilon \|\hat{\mathbf{v}}_i\|_p} \{\mathcal{L}(\hat{\mathbf{v}}_i + \Delta \mathbf{v}, \boldsymbol{\lambda}, \boldsymbol{\mu}; \mathbf{z})\} \right\} \\ &= \max_{\substack{\boldsymbol{\lambda} \geq \mathbf{0} \\ \boldsymbol{\mu} \in \mathbb{R}}} \left\{ \hat{\mathbf{v}}_i \cdot [\mathbf{z} - \boldsymbol{\lambda}] - \epsilon \|\hat{\mathbf{v}}_i\|_p \|\mathbf{z} - \boldsymbol{\lambda} + \boldsymbol{\mu} \mathbf{1}\|_p^* \right\} \end{aligned}$$

– where $\|\cdot\|_p^*$ is dual to the p -norm. The third identity allows us to represent the hypograph of u_i in the direct uncertainty model as

$$\begin{aligned} H_i &= \{(\mathbf{z}, t) \mid (\mathbf{z}, t, \boldsymbol{\mu}, \boldsymbol{\tau}, \boldsymbol{\lambda}) \in \mathbb{R}^{m+3} \times \mathbb{R}_+^m \\ &\quad \|\mathbf{z} - \boldsymbol{\lambda} + \boldsymbol{\mu} \mathbf{1}\|_p^* \leq \tau / (\epsilon \|\hat{\mathbf{v}}_i\|_p) \\ &\quad t + \tau \leq \hat{\mathbf{v}}_i \cdot [\mathbf{z} - \boldsymbol{\lambda}]\}. \end{aligned}$$

Now we address the buyerside uncertainty model. The derivation proceeds in an identical fashion, by minimizing an appropriate Lagrangian over the compact convex set $\{\boldsymbol{\theta} \mid \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_i\|_p \leq \epsilon \|\hat{\boldsymbol{\theta}}_i\|_p\}$. The outcome of this process is that $u_i(\mathbf{z}) = \min\{\mathbf{z} \cdot \mathbf{v} \mid \mathbf{v} \in \mathcal{V}_i^b(p, \epsilon)\}$ has hypograph

$$\begin{aligned} H_i &= \{(\mathbf{z}, t) \mid (\mathbf{z}, t, \boldsymbol{\mu}, \boldsymbol{\tau}, \boldsymbol{\lambda}) \in \mathbb{R}^{m+3} \times \mathbb{R}_+^m \\ &\quad \|\boldsymbol{\Phi}^\top [\mathbf{z} - \boldsymbol{\lambda} + \boldsymbol{\mu} \mathbf{1}]\|_p^* \leq \tau / (\epsilon \|\hat{\mathbf{v}}_i\|_p) \\ &\quad t + \tau \leq \hat{\mathbf{v}}_i \cdot [\mathbf{z} - \boldsymbol{\lambda}]\}. \end{aligned}$$

In the next section, we provide numerical experiments that investigate robust equilibria with respect to direct and buyerside models.

In contrast to the direct and the buyer-side uncertainty models which have tractable representations, the convex hull of the fully general joint uncertainty model given by $\mathcal{V}_i^J(\epsilon_1, \epsilon_2)$ does not seem to be amenable to an exact, tractable representation. In the special case when p corresponds to the Euclidean norm and q corresponds to the Frobenius norm, we provide a tractable representation of an *outer approximation* of this set. Robust equilibria may be computed with respect to these outer approximations; since

$$\min_{\mathbf{v} \in \mathcal{V}_i^{out}(\epsilon_1, \epsilon_2)} \mathbf{v} \cdot \mathbf{z} \leq \min_{\mathbf{v} \in \mathcal{V}_i^J(\epsilon_1, \epsilon_2)} \mathbf{v} \cdot \mathbf{z} \quad \forall \mathbf{z} \in \mathbb{R}^m,$$

the resulting allocations will be guaranteed to be robust against all uncertainty realizations in $\mathcal{V}_i^J(\epsilon_1, \epsilon_2)$.

Lemma 2. *Let $\mathcal{V}_i^{out}(\epsilon_1, \epsilon_2)$ denote the set*

$$\begin{aligned} \{\mathbf{v} : \mathbf{v} &= \boldsymbol{\mu} + \boldsymbol{\delta} \hat{\boldsymbol{\Phi}}^\top + \hat{\boldsymbol{\theta}}_i \boldsymbol{\Delta}^\top + \hat{\boldsymbol{\theta}}_i \hat{\boldsymbol{\Phi}}^\top, \\ \mathbf{v} &\geq \mathbf{0}, \mathbf{v} \cdot \mathbf{1} = \hat{\mathbf{v}}_i \cdot \mathbf{1}, \|\boldsymbol{\delta}\|_2 \leq \epsilon_1 \|\hat{\boldsymbol{\theta}}_i\|_2, \\ \|\boldsymbol{\mu}\|_2 &\leq \epsilon_1 \epsilon_2 \|\hat{\boldsymbol{\theta}}_i\|_2 \|\hat{\boldsymbol{\Phi}}\|_F, \|\boldsymbol{\Delta}\|_F \leq \epsilon_2 \|\hat{\boldsymbol{\Phi}}\|_F\}. \end{aligned}$$

Then

$$\text{conv}(\mathcal{V}_i^J(\epsilon_1, \epsilon_2)) \subseteq \mathcal{V}_i^{out}(\epsilon_1, \epsilon_2).$$

Note that $\mathcal{V}_i^{out}(\epsilon_1, \epsilon_2)$ is a convex body, and that it is tractable to optimize over since it only involves second-order cone constraints.

5 Experimental Results



We construct low rank markets following Kroer et al. (2019). We start with the MovieLens 1M dataset where ≈ 6000 individuals give ratings to ≈ 4000 movies. We use standard techniques to complete the matrix and take the 200 individuals with the most movies rated; these individuals are endowed with a unit budget of fictitious currency. From this we construct a “plentiful” market with $m = 500$ goods (movies), and a “scarce” market with $m = 50$ goods. We consider these markets in Fisher and quasi-Fisher settings, and with both direct and buyerside models from Section 4. For the buyerside models, $\hat{\boldsymbol{\Phi}}$ are of size $m \times d$, where d is the rank of the nominal valuation matrix; our scarce and plentiful markets had ranks 25 and 35 respectively. Experiments here use the 2-norm; refer to supplementary material for the same experiments under 1-norm uncertainty.

For our implementation we rely on CVXPY 1.0 (Diamond and Boyd 2016; Agrawal et al. 2018) to interface with solvers MOSEK (Mosek 2010; Dahl and Andersen 2019) and ECOS (Domahidi, Chu, and Boyd 2013). MOSEK is used to solve the equilibrium problems; although other solvers exist which support the logarithmic terms in the objective, it is our experience that no other solver can reliably handle equilibrium problems beyond very small scales. EOCS is used to evaluate utility functions as part of computing robust envy and robust-utility Nash welfare (the latter metric we define momentarily).

5.1 Nash welfare



Nash Social Welfare (the product of agent’s utility functions) is a popular measure of community utility. Here we adopt a normalized version of Nash Social Welfare: the budget-weighted geometric means of utility functions; we consider this quantity with respect to nominal utilities $\mathbf{x}_i \mapsto \hat{\mathbf{v}}_i \cdot \mathbf{x}_i + r_i$ and robust utilities $\mathbf{x}_i \mapsto u_i(\mathbf{x}_i) + r_i$. By evaluating the nominal-utility Nash welfare at the robust solution, we get a sense of the “price of robustness.” By evaluating the robust-utility Nash welfare at the nominal solution, we can measure the price of overconfidence in the point estimate $\{\hat{\mathbf{v}}_i\} \approx \mathcal{V}_i$.

In every single experiment we conducted, the nominal Nash welfare of the robust solution decayed slower than the robust Nash welfare of the nominal solution. This is to say:



the potential price of robustness was usually modest, relative to the potential price of optimism.

The robust Nash welfare of a nominal solution is very sensitive to direct uncertainty, while it is relatively stable for buyerside uncertainty. This trend holds in Fisher and quasi-Fisher markets. In quasi-Fisher models, the nominal Nash welfare of a robust solution can be *larger* than the nominal Nash welfare of the nominal solution. This is surprising, since robust solutions are optimizing for a different objective than nominal Nash welfare.



Qualitatively, this can be attributed to uncertainty causing some buyers to exit the market, which drives down prices for those that remain in the market. We found that this mostly happens with the buyerside uncertainty model in scarce markets. It is unclear why this happens more with buyerside than direct uncertainty models, however it is very reasonable that this happens more with scarce rather than plentiful markets.


5.2 Robust envy

In the last section we saw how the robust Nash welfare of a nominal solution is very sensitive to direct uncertainty, and is relatively stable for buyerside uncertainty. These trends also hold for robust envy.


The benefits of a robust solution persist in scarce markets, although the effects are less pronounced here. Out of all experiments we conducted, the figure below illustrates the case where robust envy distributions exhibited the most overlap. Even in this case, there is a clear performance benefit of the robust solution, compared to the nominal solution.

5.3 Equilibrium prices in quasi-Fisher markets

The figure below plots equilibrium prices of every good j as the uncertainty radius ϵ ranges from 0 to 0.5. These lines have reduced opacity, so that areas of higher price density can be easily discerned. The minimum, maximum, and mean prices are traced by solid black lines.



There are several trends of note in this figure. First, we see that in the plentiful market, the average price remains constant even as uncertainty increases. Thus in the plentiful setting, even very large amounts of uncertainty may not dissuade buyers from participating in the market. In the scarce market we see a different outcome: the average price drops slowly and steadily, from just below 4.0 to just above 3.5. There are still important commonalities between the plentiful and scarce markets. The simplest properties are that prices p_j do not evolve monotonically as ϵ increases, and that uncertainty can cause changes in the order of goods



when sorting by p_j . Another crucial property is that as ϵ gets particularly large, the prices converge to a common value. Such convergence agrees with our intuition that for large enough ϵ , all buyers' uncertainty sets will reduce to scaled standard simplices $\mathcal{V}_i^d(p, \epsilon) = \{v \mid 0 \leq v, v \cdot \mathbf{1} = \hat{v}_i \cdot \mathbf{1}\}$.

As a final point, we consider how different pricing schemes affect the purchasing decisions of buyers with robust utility functions. The following plot shows how adjusting prices to reflect uncertainty (in red) results in much larger revenue than if prices were set as though there was no uncertainty (in blue).

6 Acknowledgements

The majority of this research was conducted while the first and second authors were employed by Facebook, as intern and postdoc, respectively. The first author was supported (in part) by a National Science Foundation Graduate Research Fellowship.

References

Aghassi, M., and Bertsimas, D. 2006. Robust game theory. *Mathematical Programming* 107(1-2):231–273.

Agrawal, A.; Verschueren, R.; Diamond, S.; and Boyd, S. 2018. A rewriting system for convex optimization problems. *Journal of Control and Decision* 5(1):42–60.

Albert, M.; Conitzer, V.; Lopomo, G.; and Stone, P. 2017. Mechanism design with correlated valuations: Efficient methods for revenue maximization. *Under Submission at Operations Research*.

Balseiro, S. R., and Gur, Y. 2017. Learning in repeated auctions with budgets: Regret minimization and equilibrium. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17*, 609–609. New York, NY, USA: ACM.

Balseiro, S. R.; Besbes, O.; and Weintraub, G. Y. 2015. Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Science* 61(4):864–884.

Balseiro, S.; Kim, A.; Mahdian, M.; and Mirrokni, V. 2017. Budget management strategies in repeated auctions. In *Proceedings of the 26th International Conference on World Wide Web*, 15–23. International World Wide Web Conferences Steering Committee.

Ben-Tal, A., and Nemirovski, A. 2002. Robust optimization—methodology and applications. *Mathematical Programming* 92(3):453–480.

Ben-Tal, A.; El Ghaoui, L.; and Nemirovski, A. 2009. *Robust optimization*, volume 28. Princeton University Press.

Bergemann, D., and Morris, S. 2005. Robust mechanism design. *Econometrica* 73(6):1771–1813.

Bertsimas, D., and Sim, M. 2004. The price of robustness. *Operations research* 52(1):35–53.

Bertsimas, D.; Brown, D. B.; and Caramanis, C. 2011. Theory and applications of robust optimization. *SIAM review* 53(3):464–501.

Birnbaum, B.; Devanur, N. R.; and Xiao, L. 2011. Distributed algorithms via gradient descent for fisher markets. In *Proceedings of the 12th ACM conference on Electronic commerce*, 127–136. ACM.

Borgs, C.; Chayes, J.; Immorlica, N.; Jain, K.; Etesami, O.; and Mahdian, M. 2007. Dynamics of bid optimization in online advertisement auctions. In *Proceedings of the 16th international conference on World Wide Web*.

Budish, E., and Cantillon, E. 2012. The multi-unit assignment problem: Theory and evidence from course allocation at harvard. *American Economic Review* 102(5):2237–71.

Chakraborty, S.; Devanur, N. R.; and Karande, C. 2010. Market equilibrium with transaction costs. In *International Workshop on Internet and Network Economics*, 496–504. Springer.

Chen, L.; Ye, Y.; and Zhang, J. 2007. A note on equilibrium pricing as convex optimization. In *International Workshop on Web and Internet Economics*, 7–16. Springer.

Cole, R.; Devanur, N. R.; Gkatzelis, V.; Jain, K.; Mai, T.; Vazirani, V. V.; and Yazdanbod, S. 2017. Convex program duality, fisher markets, and nash social welfare. In *18th ACM Conference on Economics and Computation, EC 2017*. Association for Computing Machinery, Inc.

Conitzer, V.; Kroer, C.; Sodomka, E.; and Stier-Moses, N. E. 2018. Multiplicative pacing equilibria in auction markets. In *International Conference on Web and Internet Economics*.

Conitzer, V.; Kroer, C.; Panigrahi, D.; Schrijvers, O.; Sodomka, E.; Stier-Moses, N. E.; and Wilkens, C. 2019. Pacing equilibrium in first-price auction markets. In *Proceedings of the 2019 ACM Conference on Economics and Computation*. ACM.

Dahl, J., and Andersen, E. D. 2019. A primal-dual interior-point algorithm for nonsymmetric exponential-cone optimization.

Diamond, S., and Boyd, S. 2016. Cvxpy: A python-embedded modeling language for convex optimization. *The Journal of Machine Learning Research* 17(1):2909–2913.

Domahidi, A.; Chu, E.; and Boyd, S. 2013. ECOS: An SOCP solver for embedded systems. In *European Control Conference (ECC)*, 3071–3076.

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

Eisenberg, E. 1961. Aggregation of utility functions. *Management Science* 7(4):337–350.

Kroer, C.; Peysakhovich, A.; Sodomka, E.; and Stier-Moses, N. E. 2019. Computing large market equilibria using abstractions. *arXiv preprint arXiv:1901.06230*.

Lopomo, G.; Rigotti, L.; and Shannon, C. 2018. Uncertainty and robustness of surplus extraction. *arXiv preprint arXiv:1811.01320*.

Mosek, A. 2010. The mosek optimization software. *Online at http://www.mosek.com* 54(2-1):5.

Peysakhovich, A., and Kroer, C. 2019. Fair division without disparate impact. *arXiv preprint arXiv:1906.02775*.

Peysakhovich, A.; Kroer, C.; and Lerer, A. 2019. Robust multi-agent counterfactual prediction. *arXiv preprint arXiv:1904.02235*.

Roth, A. E. 2015. *Who Gets What? and Why: The New Economics of Matchmaking and Market Design*. Houghton Mifflin Harcourt.