

Non-Separable, Concave Utilities are Easy – in a Perfect Price Discrimination Market Model*

Vijay V. Vazirani[†]

Abstract

Recent results, establishing evidence of intractability for such restrictive utility functions as additively separable, piecewise-linear and concave, under both Fisher and Arrow-Debreu market models, have prompted the question of whether we have failed to capture some essential elements of real markets, which seem to do a good job of finding prices that maintain parity between supply and demand.

The main point of this paper is to show that even non-separable, concave utility functions can be handled efficiently in a suitably chosen, though natural, realistic and useful, market model; our model allows for perfect price discrimination. Our model supports unique equilibrium prices and satisfies both welfare theorems.

*College of Computing, Georgia Institute of Technology, Atlanta, GA 30332–0280. Email: vazirani@cc.gatech.edu. Research supported by NSF Grants CCF-0728640 and CCF-0914732, ONR Grant N000140910755, and a Google Research Grant.

[†]A preliminary version of this paper appeared in WINE 2010. The claim in that version about quasi-concave utility functions was later found to be incorrect. The slightly modified title of this version reflects this change.

1 Introduction

The celebrated Arrow-Debreu theorem [AD54], which establishes the existence of equilibria in a very general model of the economy, has been deemed to be “highly non-constructive” since it crucially uses Kakutani’s fixed point theorem; as shown by Uzawa [Uza62], the existence of general equilibrium is equivalent to fixed point theorems. The conditions imposed on utility functions of buyers in the Arrow-Debreu theorem are very weak: continuity, quasi concavity, and non-satiation.

Over the last decade, there has been a surge of interest within theoretical computer science on studying the question of efficient computability of market equilibria – not only to provide an algorithmic ratification of Adam Smith’s “Invisible hand of the market” but also because of potential applications to new markets on the Internet. This study started with highly restricted utility functions, i.e., linear [DPSV08, Jai07], before moving to more general functions. Several specialized classes of non-linear utility functions were also solved, e.g., Cobb-Douglas, CES and Leontief, see [CV04, Ye07, CV07].

We remark that convex programs have turned out to be *the tool* for efficient computation of equilibria. For several utility functions, equilibria of the corresponding markets have been captured as optimal solutions of convex programs, which have been solved by either combinatorial algorithms (see [Vaz07, Vaz12]) or convex program solvers (see [CV07]). To our knowledge, so far all markets that have yielded to efficient algorithms, either exact or approximable to any degree of accuracy, have done so via this tool.

Regarding attacking general classes of utility functions, the obvious next question was additively separable, piecewise-linear, concave utility functions (SPLC utilities); henceforth, by “additively separable” we will mean additively separable over goods, i.e., each agent has a separate utility function for each good and her total utility is additive over goods. This long-standing open problem was resolved recently for both market models. First, [CDDT09] proved PPAD-hardness for the Arrow-Debreu model under SPLC utilities. Subsequently, PPAD-hardness was established for Fisher’s model, independently and concurrently by [CT09, VY11] and [VY11]. Membership in PPAD was established for both models, first via an indirect proof in [VY11] and later by demonstrating path following algorithms in [GMSV12]. Hence, for both market models under SPLC utilities, the complexity of computing an equilibrium has been precisely pinning down.

These results dealt a serious blow to the program of algorithmically ratifying the “Invisible hand of the market” – assuming $P \neq \text{PPAD}$, these results effectively rule out the existence of efficient algorithms for almost all general and interesting classes of “traditional” market models. On the other hand, markets in the West, based on Adam Smith’s free market principle, seem to do a good job of finding prices that maintain parity between supply and demand¹. This has prompted the question of whether we have failed to capture some essential elements of real markets in our models, see [Vaz10]. Some progress has been made on this latter question: polynomial time algorithms were given for spending constraint utilities [Vaz10] and for SPLC utilities in the Fisher model, provided perfect price discrimination is introduced in the model [GV11]. Both these works deal with additively separable utility functions.

¹For example, in the West, it is hard to see a sight that was commonplace in the Soviet Union, with massive surpluses of some goods and empty shelves of others.

Clearly, for traditional market models, the gap between the “positive” algorithmic results summarized above and the generality of the Arrow-Debreu Theorem is rather large. The main point of this paper is to show that even non-separable, concave utility functions, with the additional restrictions of continuous differentiability and non-satiation, can be handled efficiently in a suitably chosen, though natural, realistic and useful, market model; our model allows for perfect price discrimination. Our work also provides insights into the widely used practice of price discrimination, and in Section 2.1, we give an application of our market model to online display advertising marketplaces.

1.1 Price discrimination and our results

Most businesses today charge different prices from different consumers for essentially the same goods or services in order to maximize their revenues. This practice, called *price discrimination*, is not only good for businesses but also customers – without it, some customers will simply not be able to avail of certain goods or services. It is not only widespread [Var85] but is also essential for the survival of certain businesses, e.g., in the airline industry.

Price discrimination is particularly important in new industries, such as telecommunications and information services and digital goods. Traditional economic analysis, which assumes decreasing returns to scale on production, recommends pricing goods at marginal cost. However, this is not relevant to several new industries, since they have very high fixed costs and low marginal costs, and hence such prices will not even recover the fixed costs; in particular, this phenomenon is commonplace in the drug industry; the cost of discovering a drug is extremely large but the marginal cost of making a pill is very small. In these situations, product differentiation and price discrimination are an important recourse. Motivated by these considerations, price discrimination has been extensively studied in economics from many different angles; see [WMT88, Var85, Var96, Sun04, EEH98, EEH94, BT04] for just a small sampling of papers on this topic.

Price discrimination was first introduced by Pigou [Pig00], who gave the notions of first, second and third degree price discrimination. *First degree price discrimination* also called *perfect price discrimination* can only be practiced by a monopoly that has complete information about the buyers’ utility functions. The price charged from a buyer is such that her *marginal willingness to pay is equal to the marginal cost of the good*, see also [Var96]. The following quote from [Pig00] (pp. 279) makes this even more clear, “A first degree would involve the charge of a different price against all the different units of commodity, in such way that the price extracted for each was equal to the demand price for it, and no consumers’ surplus was left to the buyers.”

Clearly, in these quotes, Pigou and Varian are talking about an indivisible good. Since our market model assumes divisible goods, the natural way of interpreting Pigou’s advice is to apply it to each infinitesimal quantity of good sold. Indeed, this is what we accomplish in Section 2. We do remain with the abstraction proposed by Pigou, namely that the monopolist has complete information about the buyers’ utility functions. As in [GV11], we model this situation by introducing a *middleman*, who has complete information about the buyers’ utility functions. He buys goods from sellers in the standard way, i.e., by paying the same price for each unit of good. Then, using his information about the buyers’ utility, he knows how much the buyer is willing to pay for each (infinitesimal) quantity of the good and he charges accordingly. In general, the total amount payed by buyers to the middleman exceeds the amount he payed for the goods; the difference being his

profit. The middleman makes no profit if and only if buyers' utility functions are linear.

We give a convex program that captures equilibrium for this model. Its optimal primal and dual solutions yield equilibrium allocations, under price discrimination, as well as equilibrium prices of goods. The latter are prices that the middleman pays to the sellers of goods. As detailed in the proof of Theorem 3, the exact amount the middleman needs to charge from each buyer can also be computed from the optimal primal and dual solutions. In retrospect, our convex program is a generalization of the classic Eisenberg-Gale convex program, which captures equilibrium for Fisher markets under linear utilities (see [Vaz07] for a detailed explanation). The convex program also yields particularly simple proofs of both welfare theorems for our model.

Since the monopolist needs to have complete information about the buyers' utility functions, perfect price discrimination is primarily a useful abstraction and has somewhat limited applicability. However, in Section 2.1 we do provide a setting, namely online display advertising marketplaces, which satisfies this stringent criterion to a high degree.

2 The Market Model

Our market model is based on that of Irving Fisher [BS00]. It consists of sellers of distinct, divisible goods, a middleman and buyers. Let G and B be the set of goods and buyers, respectively. Assume that $|G| = g$, $|B| = n$, the goods are numbered from 1 to g , and the buyers are numbered from 1 to n . Assume w.l.o.g. that there is 1 unit of each good. Let $m_i \in \mathbf{Q}^+$ dollars be the money of buyer i . For each buyer i we are specified a function $f_i : \mathbf{R}_+^g \rightarrow \mathbf{R}_+$ which gives the utility derived by i as a function of allocation of goods. Each f_i is differentiable, concave and satisfies non-satiation; moreover, it is polynomial time computable. We will say that f *satisfies non-satiation* if for any allocation \mathbf{y} , there is an allocation \mathbf{y}' that weakly dominates \mathbf{y} component wise and such that $f(\mathbf{y}') > f(\mathbf{y})$. Since there is one unit of each good, it is sufficient if each f_i is defined over the domain $[0, 1 + \Delta]^g$, where $\Delta > 0$ is a constant, and satisfies non-satiation for allocations in this set.

The buyers have no access to the sellers; however, the middleman does. The middleman also has full information about the utility function of each buyer. He buys goods from the sellers and sells them to the buyers. The sellers and the middleman have no utility for the goods and both are trying to maximize their earnings. The buyers have no utility for money and are interested in obtaining a utility-maximizing bundle of goods.

In this situation, perfect price discrimination is natural. Since the buyers have access to goods only through the middleman and the latter has full knowledge of their utility functions, he is able to maximize his profit by extracting all the consumer surplus.

The middleman buys goods from the seller, who charges the middleman in the usual manner, i.e., depending on the prices of goods and the amounts bought. However, in selling goods, the middleman charges buyers on the basis of the utility they accrue rather than the amount of goods they receive. The charging mechanism detailed below ensures, as promised in Section 1.1, that the rate (i.e., money charged per unit of good) at which the middleman charges a buyer for each infinitesimal quantity of a good is proportional to rate at which the buyers derives utility at this point in the allocation, with the minimum rate being equal to the actual price per unit of the good.

Henceforth, we will use the term *rate* in a different sense than that in the previous paragraph. Let r_i be the rate at which buyer i gets utility per dollar charged from her, at any given prices \mathbf{p} . It should be easy to see that the requirements of the previous paragraph can be achieved if the middleman fixes each buyer's rate appropriately, as a function of prices \mathbf{p} . Below we will show something stronger: we will show that under the following two assumptions, each buyer can be allowed to *determined her rate so as to maximize the utility she accrues*.

1. Each buyer has no utility for money but wants to maximize the utility she accrues.
2. The middleman refuses to sell any part of a good at a loss – the fact that the middleman knows buyers' utility functions enables him to do this (we will specify in Section 2.2 what this restriction means mathematically).

We show in Lemma 2 that under these assumptions, there is a unique rate r_i for each buyer i , as a function of prices. This is also the rate at which each buyer's marginal willingness to pay equals the marginal prices of goods she gets, as required under perfect price discrimination. At this rate r_i , the total utility buyer i is able to get will be $r_i \cdot m_i$. In our model, the elasticity among consumers leads to profit for the middleman; in particular, if the utility functions of all buyers are linear, then the middleman will make no profit.

We will use the following notation and definitions throughout. \mathbf{x} will denote allocations made of all goods to all buyers. \mathbf{x}_i will denote the restriction of \mathbf{x} to allocations made to buyer i only, and x_{ij} will denote the amount of good j allocated to buyer i . For the sake of ease of notation, let us introduce the following w.r.t. a generic buyer: \mathbf{y} , a vector of length g , will denote allocation and its j th component, y_j , will denote the allocation of good j . $f : \mathbf{R}_+^g \rightarrow \mathbf{R}_+$ will denote her utility function. Function f is *concave* if for any allocations \mathbf{y} and \mathbf{y}' ,

$$f\left(\frac{\mathbf{y} + \mathbf{y}'}{2}\right) \geq \frac{f(\mathbf{y}) + f(\mathbf{y}')}{2}.$$

The overall objective is to find prices for goods such that under these transactions, the market clears, i.e., there is no surplus or deficiency of any good. These will be called *equilibrium prices*. More formally, let \mathbf{p} be prices of goods and \mathbf{r} be the corresponding rates of buyers, as given by Lemma 2. Assume that each buyer i is charged at rate r_i and is allocated a bundle of goods. We will say that prices \mathbf{p} are equilibrium prices if they satisfy the following **conditions**:

1. Each good having positive price is completely sold.
2. The money spent by each buyer i equals m_i and i accrues utility at rate r_i described above and formally defined in Section 2.2.

2.1 Applying non-separable utilities to online display advertising marketplaces

[GV11] show how the perfect price discrimination market model can be applied to online display advertising marketplaces. Their model assumes additively separable utilities. We show that extending the model from separable to non-separable utilities makes it even more relevant to this application.

We first recall the setting from [GV11], which applies to companies that sell ad slots on web sites to advertisers. [GV11] view such a company as the middleman, the owners of web sites as sellers and the advertisers as buyers. They view ad slots on different web sites as goods which need to be priced. An advertiser's utility for a particular ad slot is determined by the probability that her ad will get clicked if it is shown on that slot. Advertisers typically pay at a fixed rate to the middleman for the expected number of clicks they get, i.e., they are paying at fixed rate for every unit of utility they get. Using knowledge of the utility function of buyers, the middleman is able to price discriminate.

In the model of [GV11], an advertiser's total utility is additive over all the slots she is allocated. We note that the utility to an advertiser from placing ads on multiple web sites would typically be an involved, non-separable function because the web sites may be substitutes, complements, etc. Hence, extending to non-separable utilities makes the model more relevant to this application.

2.2 Determining buyers' rates

With respect to any prices, we will give a closed-form definition of each buyer i 's rate, r_i ; for ease of notation, we will do this for the generic buyer, i.e., we will define her rate r^* . For this section, assume that prices of goods are set to \mathbf{p} . In Lemma 2 we will show r^* is indeed her optimal rate, i.e., it maximizes her utility. In Section 3 we will show that the solution of the convex (nonlinear) program will assign utilities to a buyer at precisely this rate w.r.t. equilibrium prices. Hence, there is no need to explicitly compute buyers' rates.

Given two allocations \mathbf{y} and \mathbf{y}' , we will say that \mathbf{y} *weakly dominates* \mathbf{y}' if for each good j , $y_j \geq y'_j$. The next lemma gives the mathematical condition that an allocation needs to satisfy in order to satisfy Condition 2.

Lemma 1 *An allocation \mathbf{y} made by the middleman at rate r satisfies Condition 2 iff*

$$\forall \mathbf{y}' \text{ s.t. } \mathbf{y} \text{ weakly dominates } \mathbf{y}', \forall j : \left(\frac{\partial f}{\partial y_j}(\mathbf{y}') \div r \right) \geq p_j.$$

Proof: Conceptually, assume that the middleman is making an allocation to the buyer gradually and continuously and is charging the buyer at the rate of r units of utility per dollar. Clearly, the effective price at which he is selling her good j depends on the allocation made already. If the latter is \mathbf{y} , then the *marginal price of good j at allocation \mathbf{y}* is

$$\frac{\partial f}{\partial y_j}(\mathbf{y}) \div r.$$

Therefore, at this point the middleman is selling good j at a loss iff

$$\left(\frac{\partial f}{\partial y_j}(\mathbf{y}) \div r \right) < p_j,$$

since then he is charging the buyer less for good j than the amount charged from him by the seller. The lemma follows. \square

Let us say that an allocation \mathbf{y} is *feasible for rate* r if it satisfies the condition given in Lemma 1. We next define, for each $r > 0$, the entire set of allocations that are feasible,

$$S(r) = \left\{ \mathbf{y} \mid \left(\frac{\partial f}{\partial y_j}(\mathbf{y}) \div p_j \right) \geq r \text{ if } y_j > 0 \text{ and } \leq r \text{ otherwise} \right\}.$$

The function $U : \mathbf{R}_+ \rightarrow \mathbf{R}_+$ gives the largest utility attained by a feasible allocation at rate r :

$$U(r) = \sup\{f(\mathbf{y}) \mid \mathbf{y} \in S(r)\}.$$

Clearly, if $r_1 > r_2$, $S(r_1) \subseteq S(r_2)$ and therefore $U(r_1) \leq U(r_2)$. Hence U is a decreasing function of r . Observe that because of the non-satiation condition, $\lim_{r \rightarrow 0} U(r)$ is unbounded.

Finally, we define rate r^* as follows

$$r^* = \arg \max_r \{U(r) \geq m \cdot r\},$$

where m is the money of the generic buyer. Since function $U(r)$ is unbounded as $r \rightarrow 0$, r^* is well defined for all m .

Lemma 2 r^* maximizes the utility accrued by the generic buyer.

Proof : Since $U(r)$ gives the maximum utility that the buyer can get from a feasible allocation at rate r , and since this is a decreasing function of r , if the rate is fixed at $\beta > r^*$, then the utility accrued by the buyer is $U(\beta) \leq U(r^*)$. The maximum amount of utility the buyer can get for money m at rate r is $r \cdot m$. Therefore, if the rate is fixed at $\alpha < r^*$, then the utility accrued by i is $\alpha \cdot m < r^* \cdot m$. Hence r^* is the optimal rate. \square

3 The Convex Program

Optimal primal and dual solutions of convex program (1) give equilibrium allocations and prices, respectively. Besides non-negativity, the only constraint is that at most 1 unit of each good is sold. We will denote the Lagrange variables corresponding to these constraints as p_j 's and will show that at optimality, they will be equilibrium prices of the corresponding market.

$$\begin{aligned} & \text{maximize} && \sum_{i \in B} m_i \log(f_i(\mathbf{x}_i)) && (1) \\ & \text{subject to} && \forall j \in G : \sum_{i \in B} x_{ij} \leq 1 \\ & && \forall i \in B, \forall j \in G : x_{ij} \geq 0 \end{aligned}$$

The KKT conditions for this convex program are:

1. $\forall j \in G : p_j \geq 0.$

2. $\forall j \in G: p_j > 0 \Rightarrow \sum_{i \in B} x_{ij} = 1.$
3. $\forall i \in B, \forall j \in G: p_j \geq \frac{m_i}{f_i(\mathbf{x}_i)} \cdot \frac{\partial f_i}{\partial x_{ij}}(\mathbf{x}_i).$
4. $\forall i \in B, \forall j \in G: x_{ij} > 0 \Rightarrow p_j = \frac{m_i}{f_i(\mathbf{x}_i)} \cdot \frac{\partial f_i}{\partial x_{ij}}(\mathbf{x}_i).$

Theorem 3 *The optimal primal and dual solutions to convex program (1) give equilibrium allocations and prices, and the latter are unique. Moreover, both can be computed to any required degree of accuracy in polynomial time.*

Proof : Because utility functions are assumed to be continuously differentiable, for an optimal solution to program (1) there is a unique dual, i.e., prices, that satisfies the KKT conditions stated above. From these, we will derive the 3 conditions defining equilibrium. Clearly, a separation oracle for the constraints and objective function can be implemented in polynomial time. Moreover, since the constraints are all linear, by [GLS88], the optimal solutions can be computed in polynomial time to any required degree of accuracy.

The first equilibrium condition is implied by the KKT conditions 1 and 2. Consider buyer i . Because of non-satiation, $x_{ij} > 0$ for some j . For this j , let

$$r_i = \left(\frac{\partial f_i}{\partial x_{ij}}(\mathbf{x}_i) \div p_j \right).$$

By KKT condition 4, any good j with $x_{ij} > 0$ must satisfy this equality, and if for some good j , $x_{ij} = 0$, then by KKT condition 3,

$$r_i \geq \left(\frac{\partial f_i}{\partial x_{ij}}(\mathbf{x}_i) \div p_j \right).$$

This proves that the middleman does not sell any part of a good at a loss. Substituting r_i back in KKT condition 4, we get $m_i \cdot r_i = f_i(\mathbf{x}_i)$, thereby proving that all money of buyer i is spent and r_i is the rate whose existence is established in Lemma 2. \square

4 The Welfare Theorems

The *first welfare theorem* states that allocations made at equilibrium prices are Pareto optimal and the *second welfare theorem* states that for any Pareto optimal utilities \mathbf{u}^* , there is a way of setting the initial moneys of buyers in such a way that an equilibrium obtained for this instance gives precisely \mathbf{u}^* utilities to buyers.

Theorem 4 *Both the welfare theorems are satisfied by our market model.*

Proof : If an allocation is not Pareto optimal, then it cannot maximize the objective function of program (1). Hence, the first statement follows from Theorem 3.

Let S be the image of the map from the polytope of feasible allocations to utilities accrued. Because of concavity of the utility functions, S is a convex region in \mathbf{R}_+^n . Let \mathbf{u}^* be any Pareto optimal utilities. Clearly, \mathbf{u}^* must lie on the boundary of this convex region. Let $\sum_{i \in B} a_i \cdot u_i = c$ be the hyperplane that is tangent to region S and contains the point \mathbf{u}^* . By Pareto optimality of \mathbf{u}^* , $a_i \geq 0$, for each $i \in B$, with at least one of these values being positive.

Let $m_i = a_i \cdot u_i^*$. Find the largest value of c' so that $\sum_{i \in B} m_i \log u_i = c'$ contains the point \mathbf{u}^* . Now, it is easy to see that for the chosen values of m_i 's, $\sum_{i \in B} m_i \log u_i = c'$ must be tangent to S at \mathbf{u}^* . Therefore, if the market is run with the money of each buyer i set to m_i , the equilibrium utilities will be \mathbf{u}^* , hence proving the second welfare theorem. \square

References

- [AD54] K. Arrow and G. Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22:265–290, 1954.
- [BS00] W. C. Brainard and H. E. Scarf. How to compute equilibrium prices in 1891. *Cowles Foundation Discussion Paper*, (1270), 2000.
- [BT04] V. Bhaskar and Ted To. Is perfect price discrimination really efficient? An analysis of free entry. *The RAND journal of economics*, 35(4):762–76, 2004.
- [CDDT09] X. Chen, D. Dai, Y. Du, and S.-H. Teng. Settling the complexity of Arrow-Debreu equilibria in markets with additively separable utilities. In *FOCS*, 2009.
- [CT09] X. Chen and S.-H. Teng. Spending is not easier than trading: on the computational equivalence of Fisher and Arrow-Debreu equilibria. *Journal of the ACM*, 56(3), 2009.
- [CV04] B. Codenotti and K. Varadarajan. Efficient computation of equilibrium prices for markets with Leontief utilities. In *ICALP*, 2004.
- [CV07] B. Codenotti and K. Varadarajan. Computation of market equilibria by convex programming. In N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, pages 135–158. Cambridge University Press, 2007.
- [DPSV08] N. Devanur, C. H. Papadimitriou, A. Saberi, and V. V. Vazirani. Market equilibrium via a primal-dual-type algorithm. *JACM*, 55(5), 2008.
- [EEH94] A. Edlin, M. Epelbaum, and W. Heller. Surplus maximization and price discrimination in general equilibrium: Part I. Technical Report 9405, Centro de Investigacion Economica, ITAM, 1994.
- [EEH98] A. S. Edlin, M. Epelbaum, and W. P. Heller. Is perfect price discrimination really efficient?: Welfare and existence in general equilibrium. *Econometrica*, 66(4):897–922, 1998.

- [GLS88] M. Grotschel, L. Lovasz, and A. Schirjver. *Geometric Algorithms and Combinatorial Optimization*. Springer-Verlag, 1988.
- [GMSV12] Jugal Garg, Ruta Mehta, Milind Sohoni, and Vijay V. Vazirani. A complementary pivot algorithm for market equilibrium under separable piecewise-linear concave utilities. In *ACM Symposium on the Theory of Computing*, 2012.
- [GV11] G. Goel and V. V. Vazirani. A perfect price discrimination market model with production and a rational convex program for it. *Mathematics of Operations Research*, 36:762–782, 2011.
- [Jai07] K. Jain. A polynomial time algorithm for computing the Arrow-Debreu market equilibrium for linear utilities. *SIAM Journal on Computing*, 37(1):303–318, 2007.
- [Pig00] A. C. Pigou. *The Economics of Welfare*. MacMillan and Company, 1900.
- [Sun04] Arun Sundararajan. Nonlinear pricing of information goods. *Management Science*, 50(12):1660–1673, 2004.
- [Uza62] H. Uzawa. Walras’ existence theorem and brower’s fixpoint theorem. *Econ. Stud. Quart.*, pages 59–62, 1962.
- [Var85] Hal R. Varian. Price discrimination and social welfare. *American Economic Review*, 75(4):870–75, 1985.
- [Var96] Hal Varian. Differential pricing and efficiency. *First Monday*, 1(2), 1996.
- [Vaz07] V. V. Vazirani. Combinatorial algorithms for market equilibria. In N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, pages 103–134. Cambridge University Press, 2007.
- [Vaz10] V. V. Vazirani. Spending constraint utilities, with applications to the Adwords market. *Mathematics of Operations Research*, 35(2), 2010.
- [Vaz12] V. V. Vazirani. The notion of a rational convex program, and an algorithm for the Arrow-Debreu Nash bargaining game. *Journal of the ACM*, 59(2), 2012.
- [VY11] V. V. Vazirani and M. Yannakakis. Market equilibria under separable, piecewise-linear, concave utilities. *Journal of the ACM*, 58(3), 2011.
- [WMT88] G. Norman W.B. MacLeod and J.-F. Thisse. Price discrimination and equilibrium in monopolistic competition. *International Journal of Industrial Organization*, 6(4):429–446, 1988.
- [Ye07] Y. Ye. Exchange market equilibria with Leontief’s utility: Freedom of pricing leads to rationality. *Theoretical Computer Science*, 378:134–142, 2007.