

Allocation and Pricing of Substitutable Goods: Theory and Algorithm

Huaxia Rui

Simon Business School, University of Rochester, 500 Wilson C. Blvd, Rochester, New York 14627, USA
huaxia.rui@simon.rochester.edu

De Liu

Carlson School of Management, University of Minnesota, 321 19th Avenue South, Minneapolis, Minnesota 55455, USA, deliu@umn.edu

Andrew Whinston

McCombs School of Business, The University of Texas at Austin, 2100 Speedway, Austin, Texas 78712, USA, abw@uts.cc.utexas.edu

Motivated by the thriving market of online display advertising, we study a problem of allocating numerous types of goods among many agents who have concave valuations (capturing risk aversion) and heterogeneous substitution preferences across types of goods. The goal is both to provide a theory for optimal allocation of such goods, and to offer a scalable algorithm to compute the optimal allocation and the associated price vectors. Drawing on the economic concept of Pareto optimality, we develop an equilibrium pricing theory for heterogeneous substitutable goods that parallels the pricing theory for financial assets. We then develop a fast algorithm called SIMS (standardization-and-indicator-matrix-search). Extensive numerical simulations suggest that the SIMS algorithm is very scalable and is up to three magnitudes faster than well-known alternative algorithms. Our theory and algorithm have important implications for the pricing and scheduling of online display advertisement and beyond.

Key words: substitutable goods; resource allocation; display advertising

History: Received: May 2015; Accepted: December 2016 by Subodha Kumar, after 2 revisions.

1. Introduction

Online platforms and digital markets increasingly match customers with massive number of heterogeneous goods and services. One prominent example is online display advertising, which refers broadly to text, graphical, video, or interactive advertisements that mobile and Internet users encounter when they browse non-search web pages or interact with applications. Due to increasingly sophisticated digital tracking and predictive analytics,¹ display advertisers can now distinguish audiences at a granular level, resulting in numerous audience types: one category, for instance, could be young, male, high-income adults who love video games and live in urban areas. McAfee et al. (2010) report that advertiser campaigns can have trillions of distinct audience categories to choose from, just based on demographics, geographic location, and interests-based “behavioral” attributes. Naturally, with refined audience categories, advertisers (or even campaigns) can demonstrate heterogeneous substitution preferences. For example, a video game company may value audience categories that

include young male adults regardless of their locations, while a casino may value audience categories that include adults in close vicinity regardless of their gender or age. Thus, the casino would not mind substituting impressions from young adults with those from older adults (perhaps for a lower cost). Such heterogeneous substitution preferences also exist in many other online matching markets that feature numerous differentiated products or services, such as vacation rental marketplace (e.g., Airbnb and Home-Away) crowd-sourcing labor markets (e.g., Amazon Mechanical Turk), and micro loans (e.g., LendingClub and Prosper).

While there are gains from substituting one type of goods with another, there are also preferences that could limit substitution, such as preferences for smooth consumption over time and for cross-sectional diversification. For example, the video game company may prefer that their ad impressions reach all geographical locations; the casino may prefer that impressions be evenly distributed throughout a month. We argue that such preferences can be captured by risk aversion, a concept from the utility

theory in economics. Indeed, risk aversion has been used in insurance and finance industries since very early time for similar purposes (e.g., Markowitz 1959).

To our knowledge, existing literature has not simultaneously modeled heterogeneous substitution preferences and heterogeneous risk aversion in a unified framework. Moreover, given the nature of the applications, it is critical that any new modeling approach can handle massive number of distinct good types computationally. To fill this gap, we formulate a new allocation problem that is well-motivated from the economic theory and captures heterogeneous substitution and risk-aversion preferences. We then address a formidable challenge of developing a new theory-driven algorithm that can solve the proposed allocation problem at very large scales.

Our new allocation problem allows many types of goods to be allocated among many agents, each with a concave valuation (for modeling risk aversion) and a unique substitution preference. The objective of the problem is to maximize total realized values of all agents, subject to resource availability constraints. We call such a formulation a nonlinear allocation with substitution (NAS). The solution to such problems holds implications for online display advertising and potentially many other online matching markets.

Our model and solution approaches could be useful for digital display advertising market, which is expected to reach \$32 billion in US revenue in 2016, and continues to grow rapidly at a rate of over 10% per year (EMarketer 2016). Specifically, our approaches are particularly relevant to demand-side platforms (DSPs), which buy display ads from ad exchanges, publisher networks, and other advertising properties on behalf of their member advertisers. Because a DSP can represent many advertisers, it must allocate impressions internally among member advertisers. A critical advantage of DSPs over the conventional ad agency is their allocative efficiency (Vidakovic 2013). By more efficiently allocating impressions among advertisers (or ad campaigns), DSPs can realize higher advertiser value, which in turn enables them to charge a higher fee and attract more advertisers in a long run. For this reason, this study focuses on maximizing allocative efficiency in the NAS problem. In section 6, we discuss the implications of our problem for display advertising in more detail.

Our formulation differs from most prior approaches to the advertising allocation problem in that we follow an economic approach to model advertiser preferences rather than relying on ad hoc specifications. For example, in the literature review, we contrast our approach with several existing approaches for addressing advertisers' need for

diversifying across several audience categories. While the formulation of NAS is motivated by the problem of allocating display advertising, it is well suited for allocation problems in sharing economy such as traveler-room matching in vacation rental marketplace and task allocation in crowd-sourcing labor markets (Ho and Vaughan 2012). In these markets, the number of distinctive types of tasks and services are high, and customers often have heterogeneous substitution preferences.

The contributions of this study are twofold: first, we provide a theory for allocating and pricing numerous types of goods given the heterogeneous substitution and risk-aversion preferences. The theory addresses, for example, the existence of a price vector and a corresponding allocation such that all price-taking agents find their allocation optimal for the given prices. It also provides solid foundation for the development of a fast algorithm for solving large scale NAS problems. Second, we develop a scalable algorithm for finding an optimal allocation of such goods in a time-constrained environment, which is particularly important because many NAS problems require fast computation. Our simulation results suggest that our algorithm can solve much larger problems than generic optimization algorithms, and has significant advantages over existing optimization packages in terms of speed and memory consumption.

More specifically, we have developed two key theoretical findings in this study. The first is the equivalence between Pareto optimality (PO) and the existence of a price vector, a concept closely related to competitive equilibrium prices (Gul and Stacchetti 1999). Once a price vector is given, one can easily obtain the corresponding Pareto-optimal allocation by converting multiple good types into a single standard good type (a procedure we call "standardization"), thereby dramatically reducing the dimension of the problem. Our second key theoretical insight is the finding that at least one optimal allocation is *regular*, a key new concept we introduce in response to the difficulty of directly finding the price vectors for PO allocations suggested by the first key theoretical finding. *Each regular allocation has a pseudo price vector*, one that coincides with a true price vector if the regular allocation is also PO. Unlike true price vectors, pseudo price vectors are much easier to find. More importantly, we also establish that at least one optimal allocation satisfies the regularity condition, thus we may focus only on regular allocations, which is not only *convenient* but also *sufficient*.

Based on these theoretical insights, together with a heuristic for searching the space of regular allocations indexed by indicator matrices, we develop a new algorithm called SIMS (standardization-and-indicator-matrix-search). The algorithm iterates among

regular allocation problems and solve them by the standardization technique. Our simulation results suggest that SIMS is *up to three magnitudes faster* than generic convex optimization algorithms.

It is interesting to note that many of our theoretical concepts and findings have parallels in the asset pricing theory of finance, which provides guidance on how financial assets, which yield uncertain cash flows over multiple periods, should be priced. For example, the concept of PO is closely related to the absence of arbitrage in asset pricing. Analogous to the equivalence between PO and the existence of a price vector, it is established in finance the equivalence between the absence of arbitrage and the existence of a state price vector (Ross 1978). Furthermore, our standardization technique shares the same spirit with the martingale methodology used for asset pricing (Harrison and Kreps 1979; Duffie 2001). These theoretical parallels underscore the similarity between display advertising markets and financial markets, which the literature has just begun to explore (Muthukrishnan 2009; McAfee 2011).² In this sense, our theory can be viewed as the counterpart of the asset pricing theory in the burgeoning new market for display advertising.

The SIMS algorithm we develop here is in many ways analogous to the simplex algorithm for linear programming. For example, the indicator matrices play a role as the basic solution in the simplex algorithm. The simplex algorithm iterates through basic solutions which essentially correspond to vertices of a polyhedron while the SIMS algorithm iterates through indicator matrices which essentially correspond to faces of a polyhedron. Different from the simplex algorithm, which finds the optimal solution at vertices of the polyhedron, the SIMS algorithm must go a step further to search the interior of a face of a polyhedron for an optimal solution.

We organize the rest of the study as follows: we review the related literature in section 2 and describe our research problem in section 3. In sections 4 and 5, we derive the theory and design the algorithm for NAS problem. Section 6 discusses implications of our results for online display advertising. Section 7 concludes the study.

2. Research Background

The problem of allocating heterogeneous goods among agents is a core problem of any exchange economy. Such a problem can be thought of as a transportation problem where types of good are sources and agents are destinations.³ Below, we review the connections between this research and the related transportation models and their applications to display advertising.

Our work is related to a growing display advertising literature that applies transportation models to solve the problem of allocating advertising resources. The basic problem of this literature is that given the supply of heterogeneous impressions, how to schedule the advertisements from different ad campaigns to maximize their goals. Langheinrich et al. (1999) was among the first to formulate display advertising as a linear transportation problem, where the goal is to allocate ads across different audience types to maximize the total number of estimated clicks while meeting the impression goals set by ad campaigns. Such a linear programming formulation tends to target ads on audience types where they perform the best, as measured by estimated click through rates. However, this also gives rise to an “over-targeting” problem (Chickering and Heckerman 2000, Tomlin 2000) where the optimal solution tends to show an ad to a narrow group of audience types. This is undesirable from an advertiser’s perspective, because advertisers generally prefer to spread an ad across multiple audience types (Nakamura and Abe 2005). Several subsequent studies attempt to remedy this problem by modifying the basic linear transportation problem, including imposing minimum number of impressions per audience type (Langheinrich et al. 1999; Nakamura et al., 2005) and adding a nonlinear entropy term in the objective function to force wide-spread allocation (Tomlin 2000). More recently, Turner (2012) proposed a quadratic objective function that aims to allocate impressions proportionally across all desirable audience types. The over-targeting problem reflects advertisers’ preference for diverse audience types (or “reach”), which in turn suggests there are diminishing returns associated with each audience type. Instead of heuristically patching the linear transportation model, we adopt a more theory-driven approach that directly models valuation functions with diminishing returns and the implied preference for diversity, using the utility function theory from economics. As we will illustrate, our utility function approach lends to nice economic interpretations of our findings and reveals a deep connection between the display advertising market and the financial market. Another benefit of our approach is the added flexibility of allowing heterogeneous substitution preferences across advertisers.⁴

To our knowledge, our transportation formulation has not been studied before. While our approach also results in a nonlinear (concave-valuation) transportation problem, we note that it is quite different from several other nonlinear (convex-cost) transportation problems in the literature. One type of nonlinear transportation problem, studied in the early economics literature, is the multi-facility production-transportation (P-T) problem (Sharp et al. 1970, Shetty

1959). In a P-T problem, a single type of goods is produced at and shipped from multiple plants, and the goal is to minimize total costs, which is the sum of linear transportation costs and convex production costs. Unlike the P-T formulation, we model multiple types of goods. Moreover, we also develop an algorithm to solve our problem at a very large scale.

Another related nonlinear transportation problem is the multi-commodity network flow (MCF) problem studied in the context of telecommunication networks. This literature seeks to optimally route multiple messages through a telecommunication network subject to convex congestion costs at arcs (Babonneau and Vial 2009, Ouorou et al. 2000). This literature is also concerned with solving large-scale convex MCF problems (e.g., Ouorou 2007; Babonneau and Vial 2009). Our problem differs from the convex MCF problem in at least two ways: in our problem, coupling occurs at destination nodes (via concave value functions) rather than at arcs; the MCF problem assumes identical costs for transporting messages while we allow agents to have different marginal values for goods. Due to these differences, specialized solution techniques for MCF problems cannot be used for our problem.

Our problem belongs to a class of problem called nonlinear resource allocation (NRA) problem, which, in its general form, is formulated as (see Patriksson (2008) and Katoh and Ibaraki (1998) for a review)

$$\min f(x_1, x_2, \dots, x_n), \text{ s.t. } \sum_{j=1}^n x_j = b, x_j \in [l_j, u_j], \\ \forall j = 1..n,$$

where the goal is to allocate one type of resource of a total amount b to n activities so that the objective value $f(x_1, x_2, \dots, x_n)$ is minimized. NRA problems can be classified by the type of objective functions, the type of constraints, and whether variables are integer or continuous Katoh and Ibaraki (1998). An NRA problem is said to have separable objective functions if the objective function can be written in the form of $f(x_1, x_2, \dots, x_n) = \sum_{j=1}^n f_j(x_j)$. Prior research has shown that the separable convex optimization with linear constraints is not NP-hard (Chubanov 2016, Hochbaum and Shanthikumar 1990). In contrast, non-separable NRA problems are harder, and generally have no polynomial algorithms (Hochbaum 2007). Our problem in its original form has a non-separable objective function, but can be converted into a separable NRA problem by introducing additional variables and constraints. The conversion, however, adds general linear constraints, which are not one of several special constraint types well studied in the literature (Katoh

and Ibaraki (1998)). We also note that even in the case of separable objective functions, neither of the known polynomial algorithms, except for a few quadratic optimization cases, is *strongly* polynomial (which means the running time depends on the data coefficients rather than only on the problem size) (Hochbaum 2007). The existence of strongly polynomial algorithms is still an open question.

In theory, our problem can be solved by any generic convex optimization solvers.⁵ Contemporary interior-point solvers such as LOQO (Vanderbei 1997) and MOSEK (MOSEK 2015) are generally quite effective at solving convex optimization programs with linear constraints (Bai et al. 1997; Boyd and Vandenberghe (2004)). However, when such problems have extremely high dimensions, generic convex-optimization solvers are no longer practical, as observed in the MCF literature (Ouorou et al. 2000). For applications such as display advertising, we not only need to solve extremely large problems, but also need to solve them in a timely manner, demanding specialized solution techniques for large scale problems that take advantages of the special structure of our problem formulation.

This research is broadly related to a few other literature streams on display advertising, including the ad scheduling literature and the auction literature for display advertising. The ad scheduling literature is concerned primarily with physically fitting ads into the available space and time. Though conceptually the ad scheduling problem is connected to the ad allocation problem in the sense any allocation needs to be scheduled for actual display, the ad scheduling literature has very different focus from ours. In particular, this literature focuses more on how to fit ads of different shapes into a shared space for a given audience type (Adler et al. 2002, Deane and Agarwal 2012, Kumar et al. 2006), than how to optimally match advertisements to different audience types. This literature is complementary to our study because it tends to consider more nuanced factors, such as exclusion clauses (Wilbur et al. 2013), audience externalities (Wilbur et al. 2013) and re-clicking effects (Kumar et al. 2007). A separate literature investigates the auction approach to display advertising. For example, Lahaie et al. (2008). design an auction framework that permits flexible expression of advertiser preferences. Chen et al. (2009) examine the issues of how to split the shares of impressions in a multi-winner ad auction. Liu and Viswanathan (2014) study the optimal choice of payment schedules in auctions for display advertising.

3. Problem Formulation

We assume there are M types of goods (e.g., impressions) and N agents (e.g., advertisers or ad

campaigns). We denote the set of agents by \mathcal{N} , the set of good types (*types* for short) by \mathcal{M} . Denoting the quantity of type m allocated to agent i by $x_{im} \in \mathbb{R}^0$, we formulate a nonlinear allocation problem with substitution (NAS) as follows:

$$(NAS) \max_{\{x_{im}\}} \sum_{i \in \mathcal{N}} \mathcal{U}_i(x_{i1}, x_{i2}, \dots, x_{im}) = \sum_{i \in \mathcal{N}} Q_i \left(\sum_m \alpha_{im} x_{im} \right), \quad (1)$$

$$s.t. \sum_{i \in \mathcal{N}} x_{im} \leq \omega_m, \forall m \in \mathcal{M}, \quad (2)$$

$$x_{im} \geq 0, \forall i \in \mathcal{N}, m \in \mathcal{M}, \quad (3)$$

where $\omega_m \geq 0$ represents the total supply of type m , $\mathcal{U}_i(\cdot)$ is agent i 's valuation for the portfolio $(x_{i1}, x_{i2}, \dots, x_{im})$. The objective function of the NAS problem consists of the sum of valuations $\mathcal{U}_i(\cdot)$ of all agents. The constraints (2) and (3) represent the usual feasibility and non-negativity requirements respectively. We call α_{im} agent i 's *valuation coefficient* for type m . We assume the *valuation function* $Q_i(\cdot)$ to be continuously differentiable, strictly increasing, and strictly concave for technical convenience.⁷ Our model and methods are applicable to other families of increasing and concave value functions.⁸

We say type m is *valuable* to agent i if the valuation coefficient for this type is positive ($\alpha_{im} > 0$). Because an agent's valuation has a linear core, the agent considers all valuable types as substitutes: the marginal rate of substitution is constant and is determined by the ratio of the corresponding valuation coefficients.

Our approach to modeling agent preference is rooted in the utility function theory in economics. Utility functions often take a concave form because of diminishing marginal returns. When used in a stochastic environment, a concave utility function can capture agent i 's risk aversion. Concave utility functions are widely used in insurance and finance (e.g., Markowitz 1959) and have been recently proposed as an alternative to traditional stochastic and robust programming approaches (Bai et al. 1997, Chen et al. 2007, Mulvey et al. 1995, Ye and Yao 2010).

For ease of reading, we list in Table 1 the main notations that will be used in theory development.

4. Theory

The purpose of this section is to better understand the structure of the NAS problem which will then be exploited to solve the problem efficiently, especially when the dimension is high.⁹ We prove that it is possible to break down a multi-good problem (i.e., an NAS problem with multiple types of goods, $M > 1$) to a series of much simpler single-good ones (i.e., NAS

Table 1 Summary of Notations

Symbol	Meaning
i	Index for agent
m	Index for type of good
N	Total number of agents
M	Total number of good types
\mathcal{N}	The set of agents
\mathcal{M}	The set of good types
$\mathbf{x}_i = (x_{i1}, \dots, x_{im}, \dots, x_{iM})$	Agent i 's allocation, with x_{im} being the quantity of type m goods allocated to agent i
$\mathbf{x} = (\mathbf{x}_1^T, \dots, \mathbf{x}_i^T, \dots, \mathbf{x}_N^T)^T$	The $N \times M$ allocation matrix for all agents, where \mathbf{x}_i is agent i 's allocation and T denotes matrix transpose
α_{im}	Agent i 's valuation coefficient for type m goods
$\omega = (\omega_1, \dots, \omega_m, \dots, \omega_M)$	Total supply of each type of goods, with ω_m being the total supply of type m goods
$Q_i(\cdot)$	Agent i 's valuation function
$\mathbf{p} = (p_1, p_2, \dots, p_M)$	A price vector for the M types of goods
$\tilde{\mathbf{p}} = (\tilde{p}_1, \tilde{p}_2, \dots, \tilde{p}_M)$	A pseudo price vector for the M types of goods
$\mathbf{z} = (z_1, z_2, \dots, z_N)$	The allocation of standardized goods, with z_i being the allocation of standardized goods for agent i
\mathbf{l}	Indicator matrix of dimension $N \times M$ whose element $l_{im} \in \{0, 1\}$ represents whether agent i 's is allowed to have type m goods
$\tilde{\omega}$	The standardized supply
$\tilde{Q}_i(\cdot)$	Agent i 's valuation function in terms of standardized goods
λ_m	Lagrange multiplier of the supply constraint of type m good (i.e., its shadow price)
v_i	A scalar used in Example 1 as a parameter of the $Q_i(\cdot)$ function

problems with only one type of good, $M = 1$), thus providing a foundation for an efficient iterative algorithm.

The basic idea of our theoretical analysis is to take advantage of the correspondence between *Pareto optimality* (PO)¹⁰, a necessary condition for optimality, and the existence of a price vector, under which the PO allocation is optimal for each agent (Theorem 1). We further show that given a price vector, we can reduce a multi-good NAS problem into a single-good one, a technique which we call *standardization* (Theorem 2). Finally, (Theorem 3) we establish that, in order to find the optimal PO allocation, it is sufficient to search among *regular* allocations (Theorem 3): finding a regular allocation is much easier than finding a PO allocation, and any regular allocation has a pseudo price vector which also allows the standardization procedure. These results pave the way for an efficient algorithm that iteratively searches among *regular* allocations and solve them efficiently using the standardization technique.

To our knowledge, no prior work in the transportation literature has established similar theoretical results for their models. However, as we will discuss, there are interesting analogies between Theorem 1 and the first fundamental theorem of asset pricing, between Theorem 2 and the Martingale method widely used for financial asset pricing, and between Theorem 3 and the fundamental theorem of linear programming that gives rise to the celebrated simplex method.

We develop these theoretical results in five subsections, starting from the concept of an indicator matrix which we use to denote a family of allocations. All proofs are available in Appendix A.

4.1. Graph Representation and Indicator Matrix

As with other transportation problems, our problem can also be represented in a graph where source nodes are good types, destination nodes are agents, and there is an arc connecting every source and destination pair. Thus, our problem is also a network flow problem with the goal of figuring out the optimal flow on each arc.

Instead of allowing flows from every source to every destination, it is useful to study a restricted problem where only a subset of flows are permitted. The permitted flows can be represented by an $N \times M$ indicator matrix \mathbf{I} , whose element $I_{im} \in \{0, 1\}$ represents whether a flow is allowed from source (type) i to destination (agent) m , that is:

$$I_{im} = 0 \Rightarrow x_{im} = 0, \forall i, m.$$

We can define a NAS problem restricted by indicator matrix \mathbf{I} as follows:

$$(RNAS) \max_{\{x_{im}\}} \sum_{i \in \mathcal{N}} Q_i \left(\sum_m \alpha_{im} x_{im} \right) \quad (5)$$

$$s.t. \sum_{i \in \mathcal{N}} x_{im} \leq \omega_m, \forall m \in \mathcal{M}$$

$$x_{im} \geq 0, \forall i \in \mathcal{N}, m \in \mathcal{M}, \quad (6)$$

$$x_{im} = 0, \forall i \in \mathcal{N}, m \in \mathcal{M}, I_{im} = 0. \quad (7)$$

The last condition requires the allocation matrix \mathbf{x} to have positive values only at places where the indicator matrix \mathbf{I} has “1”. The three conditions collectively define the set of *feasible allocations* for the restricted problem.

We use the following example throughout the study:

EXAMPLE 1. Consider the following 4 9 4 example with supply vector $\omega = (12, 8, 6, 6)$ and exponential valuation functions¹¹

$$Q_i(\mathbf{x}_i) = v_i \left(1 - e^{-\sum_{m=1}^4 \alpha_{im} x_{im}} \right), \quad i = 1, 2, 3, 4,$$

where the parameters v_i and the valuation coefficients α_{im} are given by:

$$\mathbf{v} = \begin{pmatrix} 2 \\ 1 \\ 1.5 \\ 1.2 \end{pmatrix}, \quad \alpha = \begin{bmatrix} 0.3 & 0.16 & 0.1 & 0.2 \\ 0.2 & 0.5 & 0.12 & 0.05 \\ 0.13 & 0.1 & 0.4 & 0.08 \\ 0.06 & 0.1 & 0.2 & 0.3 \end{bmatrix}.$$

We consider three indicator matrices for this problem:

$$\mathbf{I}^* = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{bmatrix}, \quad \mathbf{I}^1 = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{bmatrix},$$

$$\mathbf{I}^2 = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 \end{bmatrix}.$$

An NAS problem restricted by \mathbf{I}^* would allow agents 1–4 to have types $\{1\}$, $\{2\}$, $\{1, 3\}$, and $\{2, 4\}$ respectively. \mathbf{I}^1 additionally allows agent 1 to have type 2. \mathbf{I}^2 additionally allows agent 3 to have type 4.

We are interested in restricted problems that contain the solution to the original NAS problem. We call such an indicator matrix an *optimal indicator matrix*. By definition, an indicator matrix of all 1’s is always optimal. We are interested in non-trivial optimal indicator matrices with fewer 1’s.

4.2. Pareto Optimality and Price Vector

Apparently, for an allocation to be optimal, it is necessarily *Pareto optimal* (PO), which means that one cannot make some agents better off without hurting others through a reallocation of goods (i.e., no Pareto improvement). A formal definition of PO is given in Appendix A.3. We say an indicator matrix \mathbf{I} is PO if all feasible allocations in the NAS problem restricted by \mathbf{I} are PO.

The second welfare theorem of economics establishes that there is a correspondence between PO allocation and the existence of a set of competitive equilibrium prices such that all price-taking agents would prefer this allocation to any other affordable allocation. We next show that a similar insight holds for our problem. We first introduce the concept of a *price vector* and then show that the existence of a price vector is equivalent to PO.

DEFINITION 1. (*Price Vector*) A strictly positive vector $\mathbf{p} = (p_1, p_2, \dots, p_M)$ is called a price vector for an

NAS problem restricted by) an indicator matrix \mathbf{I} if

$$\frac{\alpha_{im}}{\alpha_{in}} \geq \frac{p_m}{p_n}, \forall i \in \mathcal{N}, m, n \in \mathcal{M}, \text{ such that } I_{im} = 1. \quad (8)$$

The price vector captures the idea of “equilibrium” prices in a competitive market such that if the goods were to be traded at these prices, no agent would find it profitable to do so. Condition (8) says that agent i can have type m ($I_{im} = 1$) only if her valuation for type m relative to any other type (α_{im}/α_{in}) is at least as high as the price for type m relative to any other type (p_m/p_n). In other words, if there were a decentralized market where the posted prices were \mathbf{p} , the agent would not gain by trading her current allocation for another.

Because re-scaling of \mathbf{p} would not affect condition (8), the price vector as defined above, if it exists, is clearly not unique. From now on, we say that a restricted NAS problem has a *unique* price vector if all of its price vectors are proportional to each other.

The following example shows that a price vector may not exist or be unique for an arbitrary restricted problem.

EXAMPLE 2. Continuing with Example 1, it can be verified that, in the case of \mathbf{I}^* , any vector $\mathbf{p} = (13, b, 40, 3b)$ with $10 \leq b \leq 20$ satisfies Equation (8). Hence, the price vector for \mathbf{I}^* is not unique. \mathbf{I}^1 has a unique price vector $\mathbf{p} = (3.9, 2.08, 12, 6.24)$. One can also prove that, in the case of \mathbf{I}^2 , condition (8) cannot be met, so there is no price vector for \mathbf{I}^2 .¹²

Theorem 1 below establishes the correspondence between PO and the existence of a price vector.

THEOREM 1. An indicator matrix \mathbf{I} is Pareto optimal if and only if there exists a price vector for \mathbf{I} .

We briefly explain the intuition for the proof here. We first establish that PO is equivalent to the absence of any “profitable” trading cycle where each person in a circle gives one type of her goods to the next person. In the simplest setting with two agents, 1 and 2, and two types of goods, A and B, any exchange is a trading cycle: for example, agent 1 may exchange 1 unit of type A with agent 2 for x units of type B. The existence of a price vector (plus the fact that agents 1 has A and agent 2 has B) implies that $\frac{\alpha_{1A}}{\alpha_{1B}} \geq \frac{p_A}{p_B} \geq \frac{\alpha_{2A}}{\alpha_{2B}}$. So if agent 1 finds the exchange profitable (which requires $x > \frac{\alpha_{1A}}{\alpha_{1B}}$), then agent 2 must not find it profitable (which requires $x < \frac{\alpha_{2A}}{\alpha_{2B}}$), and vice versa. Hence, there cannot be a Pareto improvement trading cycle in this setting. Conversely, if the allocation is PO, we infer that $\frac{\alpha_{1A}}{\alpha_{1A}} \geq \frac{\alpha_{2A}}{\alpha_{2B}}$, thus we can always find a price

vector that satisfies the condition $\frac{\alpha_{1A}}{\alpha_{1B}} \geq \frac{p_A}{p_B} \geq \frac{\alpha_{2A}}{\alpha_{2B}}$. Our proof generalizes the basic idea in this simple case to any number of agents and any number of good types.

It is interesting to note that Theorem 1 has a counterpart in the asset pricing theory, namely, the first fundamental theorem of asset pricing which states that a financial market is free of arbitrage if and only if there exists a state-price vector. The analogy has its root in the connection between PO and absence of arbitrage.

4.3. Price Vector and Standardization

Knowing the price vector is extremely valuable because it allows us to convert multiple types into a standard type, thus dramatically reducing the dimension of the problem, as we show in the next Theorem.

THEOREM 2. (Standardization) Let \mathbf{I} be a Pareto-optimal indicator matrix and \mathbf{p} be an associated price vector. Define the supply $\tilde{\omega}$ and valuation functions $\tilde{Q}_i(\cdot)$, $i \in \mathcal{N}$, for the “standard” good as:¹³

$$\tilde{\omega} \equiv \sum_{m \in \mathcal{M}} \omega_m p_m, \quad (9)$$

$$\tilde{Q}_i(z_i) \equiv \begin{cases} Q_i\left(\frac{z_i}{p_m} z_i\right), & \forall i \text{ such that } I_{im} = 1 \text{ for some } m \\ 0, & \text{otherwise} \end{cases} \quad (10)$$

Let \mathbf{z}^* be the solution to the following standardized single-good NAS problem

$$\begin{aligned} (\text{single-type NAS}) \quad & \max_{\{z_i\}} \sum_{i \in \mathcal{N}} \tilde{Q}_i(z_i) \\ & \text{s.t.} \quad \sum_{i \in \mathcal{N}} z_i \leq \tilde{\omega} \\ & z_i \geq 0, \forall i \in \mathcal{N} \end{aligned} \quad (11)$$

and \mathbf{x} be an allocation restricted by \mathbf{I} that satisfies the following system of linear equations:¹⁴

$$\begin{cases} \sum_{m \in \mathcal{M}, I_{im}=1} p_m x_{im} = z_i^*, \forall i \in \mathcal{N} \\ \sum_{i \in \mathcal{N}, I_{im}=1} x_{im} = \omega_m, \forall m \in \mathcal{M} \end{cases} \quad (12)$$

The allocation \mathbf{x} is a solution to the original NAS problem if it is non-negative.

Theorem 2 suggests that given a price vector, we can convert multiple good types into a standard good type. In this standardized economy, the total supply of standard goods is the sum of all original goods weighted by their prices (Equation (9)) and the valuation coefficients of standard goods are valuation coefficients of the goods divided by their prices (Equation (10)). The system of linear equations (12) allows us to recover an allocation of goods from an optimal

allocation of standard goods. More importantly, if both \mathbf{I} and the associated price vector are chosen “correctly,” the allocation \mathbf{x} recovered from the standardized problem is a solution to the original NAS problem.

In the proof of Theorem 2 (Appendix A.5), we show that if a PO indicator matrix \mathbf{I} and the associated price vector \mathbf{p} are “correct”, then \mathbf{p} is proportional to the competitive equilibrium prices¹⁵, and the scaling factor is exactly the Lagrange multiplier for the supply constraint in the standardized problem. Thus, an appropriately scaled price vector for an optimal indicator matrix can also be interpreted as the equilibrium prices in a competitive market.

It is also interesting to notice the connection between our standardization technique and martingale pricing method which has become the workhorse in the financial industry over the last few decades. To see this, we need to interpret the space of impression types as the sample space (Ω) in a probability space and the supply of numerous types of impressions as an asset with uncertain values depending on the realized outcome in the sample space. The price vector in our standardization theorem, once normalized, essentially defines a martingale probability measure (\mathbb{P}) under which the “standardized” supply is the *expected* supply ($\tilde{\omega} = \mathbb{E}_{\mathbb{P}}[\omega]$). More importantly, under this probability measure \mathbb{P} , the value of a portfolio is completely determined by its expectation under \mathbb{P} and agents care only about their *expected* allocations. Hence, we only need to allocate goods among agents based on *the expected* supply and later constructs the actual allocation that is consistent with the expected values and the supply constraints by solving a system of linear equations.

Being able to reduce a multi-good problem to a single-good one is a significant advantage, especially for a large-scale problem with numerous types of goods. Theorem 2 suggests the following iterative procedure for solving an NAS problem.

- First, we identify a PO indicator matrix \mathbf{I} and obtain a price vector.
- Second, we use the price vector to standardize the goods according to Equation (9).
- Third, we solve the standardized single-good NAS problem, which can be done relatively easily.
- Fourth, we obtain a candidate allocation for the original problem by solving the system of linear Equations (12).
- Finally, we test the optimality of the candidate allocation and if it is not optimal, we find another Pareto-optimal indicator matrix and start from step 1.

However, there are still several practical challenges. First, it is unclear how to find the first Pareto-optimal indicator matrix and, if the current one does not produce the solution to the NAS problem, how to find the next one. Though we have provided a condition for PO in Lemma 3 of Appendix A.3, directly verifying PO is far from trivial. Second, deriving a price vector from a known Pareto-optimal indicator matrix is not straightforward either, even for simple cases such as Example 2. We address these challenges in two steps: first, we introduce a new concept called *regularity*, which overlaps with PO but is much more computation-friendly; second, we introduce a heuristic matrix search algorithm in section 5 for navigating in the space of regular indicator matrices. The regularity condition is built upon the notion of *connectivity* between good types, which we discuss before the concept of regularity.

4.4. Connectivity between Types of Goods

As we have mentioned before, an indicator matrix can be alternatively thought of as describing a network of agents and good types. Types of goods are indirectly connected by agents who are linked to them. Using this notion of connectivity, we can define a connected indicator matrix.

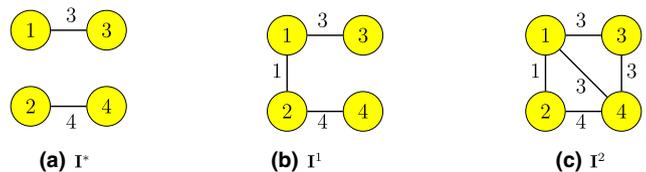
DEFINITION 2. (Connected Types) In an indicator matrix \mathbf{I} , types m and n are connected via agent i , denoted as $m \overset{i}{\leftrightarrow} n$, if the agent can have both m and n , i.e., $I_{im} = I_{in} = 1$.

Based on this notion of connectivity, we can define a *graph* G for each indicator matrix \mathbf{I} using types as nodes and connecting agents as labels. Figure 1 illustrates the connectivity graphs associated with \mathbf{I}^* , \mathbf{I}^1 , and \mathbf{I}^2 respectively.

DEFINITION 3. (Connected Indicator Matrix) An indicator matrix \mathbf{I} is connected if its graph is connected.

In Example 1, \mathbf{I}^1 and \mathbf{I}^2 are connected but \mathbf{I}^* is not. When an indicator matrix \mathbf{I} is disconnected, its graph can be decomposed into several connected components. We call the set of nodes in each connected

Figure 1 Connectivity Graphs Corresponding to \mathbf{I}^* , \mathbf{I}^1 and \mathbf{I}^2 [Color figure can be viewed at wileyonlinelibrary.com]



component of the graph a *connected component* of the indicator matrix \mathbf{I} .

Recall that a Pareto-optimal indicator matrix must have a price vector, but as we state earlier, the price vector needs not be unique. It turns out that when a Pareto-optimal indicator matrix is connected, the price vector must be unique (i.e., up to a scaling factor).

PROPOSITION 1. *If a connected indicator matrix \mathbf{I} is Pareto optimal, then the price vector for \mathbf{I} is unique.*

The intuition for this result is as follows. Whenever an agent owns two types of goods, the price ratio between these goods will be determined by the agent’s marginal valuations for them. A connected indicator matrix implies that all goods types are directly or indirectly connected, and therefore their price ratio are also determined.

As illustrated in Example 2, with each component of \mathbf{I}^* having its own price vector and scaling factor (i.e., 1 and b respectively for the two components in the example) at the component level, the price vector for the entire indicator matrix \mathbf{I}^* becomes non-unique.

4.5. Regularity

Recall that when two good types are connected by an agent, their price ratio is determined by the marginal valuations of that agent. What if two good types are connected by multiple agents? It turns out that it implies either sub-optimality or alternate solutions. Regularity rules out such conditions, and yields enormous benefits for computation.

To motivate the concept of regularity, we first consider a simple example

EXAMPLE 3. *Consider an example with two agents and two types of goods. Let*

$$U_1(\mathbf{x}_1) = Q_1(x_{11} + x_{12}), U_2(\mathbf{x}_2) = Q_2(x_{21} + \beta x_{22}).$$

Consider five connected indicator matrices

$$\mathbf{I}^a = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}, \mathbf{I}^b = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix}, \mathbf{I}^c = \begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix},$$

$$\mathbf{I}^d = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \mathbf{I}^e = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}.$$

Based on valuation coefficients, agent 1 is indifferent between the two types. Depending on the value of β , agent 2 may prefer one type to the other type or be indifferent between them.

- If $\beta < 1$, agent 2 prefers type 1 to type 2. Hence, agent 2 would be better off trading type-2 good for

type-1 good with agent 1 until agent 2 runs out of type 2 (corresponds to \mathbf{I}^b) or agent 1 runs out of type 1 (corresponds to \mathbf{I}^c). Since agent 1 is not worse off from this trade, \mathbf{I}^a is Pareto dominated by \mathbf{I}^b or \mathbf{I}^c

- If $\beta > 1$, by symmetry, \mathbf{I}^a is Pareto dominated by \mathbf{I}^d or \mathbf{I}^e .
- If $\beta = 1$, both agents are indifferent between the two types, so we can let one agent trades one of her types for another until one of the agents runs out one good type (corresponds to $\mathbf{I}^b, \mathbf{I}^c, \mathbf{I}^d$, or \mathbf{I}^e), without affecting any agent’s valuation. In other words, \mathbf{I}^a is redundant for the purpose of finding an optimal indicator matrix.

Therefore, regardless of the value of β , excluding \mathbf{I}^a , does not sacrifice optimality: for the purpose of finding optimal allocations, we can focus on \mathbf{I}^b through \mathbf{I}^e . We note that in \mathbf{I}^a , the two types of goods are connected by *two* different agents, whereas in \mathbf{I}^b through \mathbf{I}^e , each is connected by a single agent. We generalize this important insight by introducing the concept of regularity in the following steps.

DEFINITION 4. (Regular Connection) *Given an indicator matrix \mathbf{I} , a type m has a regular connection with a connected component S ($m \notin S$) if (a) m is connected to at least one element of S and (b) all of m ’s connections to S are via the same agent.*

This generalizes the notion of “connected by a single agent” to one type against a component of types.

DEFINITION 5. (Regular Connected Component) *A connected component is regular, if each type has a regular connection with each of the connected components formed by the remaining types in this component, after removal of this type.*

DEFINITION 6. (Regular Indicator Matrix) *An indicator matrix \mathbf{I} is regular if all of its connected components are regular.*

By definition, to check for regularity, we only need to ensure the regularity of each connected component of an indicator matrix. For a connected component to be regular, each type in the component must connect to each connected component of the remaining types via a single agent (but connections to different components of the remaining types need not be through the same agent). In the examples in Figure 1, \mathbf{I}^* and \mathbf{I}^1 are regular but \mathbf{I}^2 is not because, for instance, type 4 is connected to component $\{1, 2, 3\}$ via both agent 3 and agent 4.

We next show that a regular and connected indicator matrix produces a pseudo price vector that is closely related to the true price vector.

PROPOSITION 2. *Let \mathbf{I} be a connected and regular indicator matrix. Then (a) there exists a vector $\tilde{\mathbf{p}} = (\tilde{p}_1, \tilde{p}_2, \dots, \tilde{p}_M)$, called a pseudo price vector, such that for any two connected types $m \leftrightarrow n$,*

$$\frac{\tilde{p}_m}{\tilde{p}_n} = \frac{\alpha_{im}}{\alpha_{in}}. \quad (13)$$

(b) *The pseudo price vector is unique (in the same sense as a “unique” price vector). (c) If \mathbf{I} is also Pareto optimal, then the pseudo price vector is the unique price vector for \mathbf{I} .*

The results in Proposition 2 is quite intuitive. Because the regularity and connectivity conditions ensure that any pair of goods types are connected via a single chain of agents, the pseudo price vector as determined by connecting agents’ marginal valuations is unique.

The vector as defined by Equation (13) is “pseudo” because the regularity connection only speaks about the connectivity, not whether there can be Pareto improvement among connected agents. Proposition 2 suggests that the pseudo price ratio becomes a true price vector (and a unique one) when \mathbf{I} is not only connected and regular, but also PO.

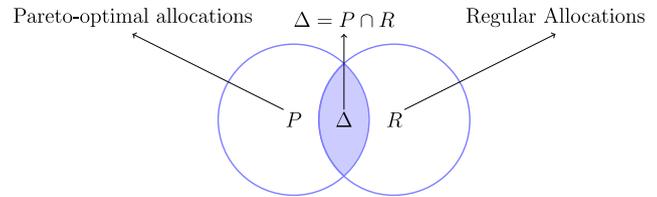
The pseudo price vector derived from condition (13) is extremely easy to compute and is a natural candidate for the price vector. To focus the search among regular indicator matrices, we must ensure that an optimal allocation resides among regular allocations. Our next result guarantees this.

THEOREM 3. (Regularity) *If a Pareto-optimal allocation x is not regular, then there exists a regular Pareto-optimal allocation x' such that all agents are indifferent between x and x' .*

The intuition behind this important theorem can be seen from Example 3. The basic idea is that if a Pareto-optimal allocation allows multiple connecting agents (thus not regular), we can initiate exchanges among these agents without hurting any agent until some agents run out of their allocated goods. This can go on until we reach a regular and still PO allocation.

Since each Pareto-optimal allocation must have an equivalent regular allocation (Theorem 3), it is sufficient to search among regular indicator matrices. Figure 2 illustrates the relations among three key concepts in this section: optimality, Pareto optimality, and regularity.

Figure 2 P is the Set of Allocations that are PO and R is the Set of Allocations that are Regular. At Least One Optimal Allocation Resides in $\Delta = P \cap R$ [Color figure can be viewed at wileyonlinelibrary.com]



It is interesting to note that Theorem 3 plays a similar role in solving NAS as the fundamental theorem of linear programming does in solving linear programming problems. The fundamental theorem of linear programming guarantees the existence of a basic optimal solution, if an optimal solution exists. Analogously, Theorem 3 ensures that there must exist a regular optimal allocation.

The following result further shows the practical importance of the concept of regularity for the algorithm design. The proof is available in Appendix A.6.

PROPOSITION 3. *If the indicator matrix \mathbf{I} is regular, then there exists a unique solution to the system of linear equations defined by Equation (12), where the price vector \mathbf{p} is replaced by $\tilde{\mathbf{p}}$, a pseudo price vector for \mathbf{I} .*

5. The SIMS Algorithm

Based on the theoretical results in section 4, we develop the SIMS (standardization-and-indicator-matrix-search) algorithm which has two major components: the standardization component that solves an RNAS problem given a regular indicator matrix, and the indicator-matrix-search component that suggests an alternative regular indicator matrix if the current one turns out to be not optimal. We elaborate the standardization procedure below and describe the indicator-matrix-search component in Appendix B.

5.1. The Standardization Procedure

Given our results on regularity (Theorem 3), the five-step procedure suggested by Theorem 2 can be implemented, using regular indicator matrices instead.

Results on connectivity and regularity suggest we can decompose a regular indicator matrix into connected components. Suppose a regular indicator matrix \mathbf{I} has J components. We denote \mathcal{M}_j as the set of types of goods within the j th component, \mathcal{N}_j as the set of affiliated agents (i.e., who are allowed to have at least one type in \mathcal{M}_j), and \mathbf{I}_j as the submatrix of \mathbf{I} corresponding to the j th component. We define a

sub-problem as allocating goods of types in \mathcal{M}_j among agents in \mathcal{N}_j restricted by indicator matrix \mathbf{I}_j .

Because the indicator matrix \mathbf{I}_j for each sub-problem is connected and regular, we can calculate the pseudo price vector and use that in place of the price vector in the standardization procedure. Once we have the solutions of all sub-problems, we have a candidate solution to the full original NAS problem, because the allocation for each good type (or agent) is determined by the sub-problem where it belongs. To illustrate this, we continue with Example 1 and solve the NAS problem restricted by \mathbf{I}^* .

EXAMPLE 4. Continue with Example 1 restricted by \mathbf{I}^* . We can rearrange the rows (agents) and columns (goods) of \mathbf{I}^* as

$$\begin{matrix} & 1 & 2 & 3 & 4 \\ \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \end{matrix} & \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 1 \end{pmatrix} \end{matrix} \quad (14)$$

With the rearrangement, it becomes clear the matrix has two disconnected components: the first component consists of agents $\{1, 3\}$ and types $\{1, 3\}$ and the second one consists of agents $\{2, 4\}$ and types $\{2, 4\}$. As the first step, we decompose \mathbf{I}^* into two sub-matrices \mathbf{I}_1^* (the top-left component in Equation (14)) and \mathbf{I}_2^* (the bottom-right component). As the second step, we standardize each sub-problem. Take \mathbf{I}_2^* as an example. Noting that types 2 and 4 are connected via agent 4, we calculate the pseudo price vector as $(p_2, p_4) = (1, 3)$ because $\alpha_{42}/\alpha_{44} = 1/3$. The standardized supply is $\tilde{\omega} = 8p_2 + 6p_4 = 26$ and the standardized valuation functions are

$$\tilde{Q}_2(z_2) = 1 - e^{-0.5z_2}, \quad \tilde{Q}_4(z_4) = 1.2 (1 - e^{-0.1z_4}).$$

The optimal solution for the standardized problem is $z_2^* = 6.7119$, $z_4^* = 19.2881$. By Equation (12), we solve the following linear equations

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 3 \\ 1 & 1 & 0 \end{pmatrix} \begin{pmatrix} x_{22} \\ x_{42} \\ x_{44} \end{pmatrix} = \begin{pmatrix} 6.7119 \\ 19.2881 \\ 8 \end{pmatrix}$$

and obtain a solution to the second sub-problem

$$\begin{pmatrix} x_{22} & x_{24} \\ x_{42} & x_{44} \end{pmatrix} = \begin{pmatrix} 6.7119 & 0 \\ 1.2881 & 6 \end{pmatrix}.$$

A similar procedure yields a solution to the first sub-problem

$$\begin{pmatrix} x_{11} & x_{13} \\ x_{31} & x_{33} \end{pmatrix} = \begin{pmatrix} 11.823 & 0 \\ 0.177 & 6 \end{pmatrix}.$$

Putting together the solutions to the two sub-problems, we obtain the following candidate solution,

$$\mathbf{x} = \begin{bmatrix} 11.823 & 0 & 0 & 0 \\ 0 & 6.7119 & 0 & 0 \\ 0.177 & 0 & 6 & 0 \\ 0 & 1.2881 & 0 & 6 \end{bmatrix}.$$

Given a candidate solution, it is straightforward to check its Optimality, using the following result.

PROPOSITION 4. Let \mathbf{x} be the candidate solution assembled from the solutions to the J sub-problems and λ_m be the Lagrange multiplier (or shadow price) for type m . \mathbf{x} is the solution to the NAS problem if \mathbf{x} is non-negative and

$$\text{premium}_{im} \equiv \frac{\partial Q_i}{\partial x_{im}} / \lambda_m - 1 \leq 0, \forall m \in \mathcal{M}_j, i \notin \mathcal{N}_j, \quad (15)$$

where premium_{im} is termed as the value premium of agent i for goods m .

Intuitively, condition (15) ensures that an agent would not prefer goods from a different component. The value premium captures the extent to which an agent values a type m above its shadow price λ_m . At an optimal allocation, no agent should have a positive value premium for any type, particularly for types from a different component. This makes intuitive sense because otherwise, we should allocate more to this user until her marginal valuation equals the shadow price.

To check whether the candidate solution from Example 4 is optimal, we compute the marginal valuation matrix $Q'(\mathbf{x})$ as

$$Q'(\mathbf{x}) = \begin{bmatrix} \frac{\partial Q_1}{\partial x_{11}} & \frac{\partial Q_1}{\partial x_{12}} & \frac{\partial Q_1}{\partial x_{13}} & \frac{\partial Q_1}{\partial x_{14}} \\ \frac{\partial Q_2}{\partial x_{21}} & \frac{\partial Q_2}{\partial x_{22}} & \frac{\partial Q_2}{\partial x_{23}} & \frac{\partial Q_2}{\partial x_{24}} \\ \frac{\partial Q_3}{\partial x_{31}} & \frac{\partial Q_3}{\partial x_{32}} & \frac{\partial Q_3}{\partial x_{33}} & \frac{\partial Q_3}{\partial x_{34}} \\ \frac{\partial Q_4}{\partial x_{41}} & \frac{\partial Q_4}{\partial x_{42}} & \frac{\partial Q_4}{\partial x_{43}} & \frac{\partial Q_4}{\partial x_{44}} \end{bmatrix} = \begin{bmatrix} \mathbf{0.017288} & 0.0092202 & 0.0057626 & 0.011525 \\ 0.0069754 & \mathbf{0.017438} & 0.0041852 & 0.0017438 \\ \mathbf{0.017288} & 0.013298 & \mathbf{0.053193} & 0.010639 \\ 0.010463 & \mathbf{0.017438} & 0.034877 & \mathbf{0.052315} \end{bmatrix},$$

where the bold-faced elements are Lagrange multipliers λ . Noticing that \mathbf{x} is non-negative and there is no positive value premium, we conclude that \mathbf{x} is an optimal allocation.

5.2. Numerical Studies

We use three sets of numerical simulations to study the performance of the SIMS algorithm. In the first set

of simulations, we are mainly interested in the convergence behavior and scalability of SIMS. In the second set of simulations, we compare SIMS to generic convex optimization solvers. In the third set, we conduct a more realistic simulation of display advertising problem, and demonstrates the applicability of SIMS for this problem.

5.2.1. Performance and Scalability. In the first set of simulations, each agent's valuation function takes an exponential form which we have used in earlier examples. The coefficient v_i are independently drawn from the uniform distribution in the interval [1000, 10,000] and the coefficients of α_{ij} are independently drawn from the uniform distribution in the interval [0.1, 1.1]. The supply of each type of good is randomly generated according to a binomial distribution with 10 trials and a success probability of 0.4.

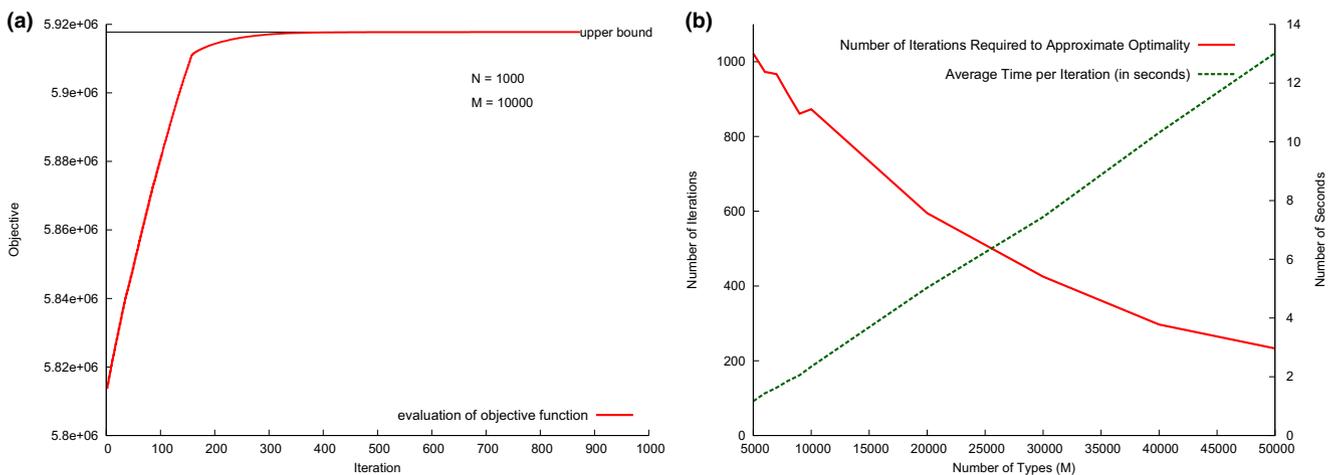
We first show that SIMS can effectively solve large-scale NAS problems and demonstrate its fast convergence. Thanks to the form of valuation functions, we can obtain a strict upper bound for the objective¹⁶, which is useful for studying the convergence behavior of SIMS. We fix the number of agents to $N = 1000$ and gradually increase the number of types from $M = 5000$ to $M = 50,000$. Because an increase in M naturally makes the allocation problem "easier" to solve due to the increase of supply, we scale down the supply vector proportionally as we scale up the value of M . This makes the convergence processes corresponding to different values of M more comparable. For all these examples, the strict upper bounds of objective values are in the interval of [5,824,870.03, 6,180,232.48]. Figure 3 plots the simulation results. The plot on the left shows the objective value at each iteration for $M = 10,000$, which quickly approaches the upper bound. This suggests that SIMS can find an approximately optimal allocation within a few

hundred iterations, which is highly valuable for practical purposes. The plot on the right further characterizes the convergence behavior of SIMS in terms of the number of iterations it takes to converge to 99.999999% of the upper bound of the objective, and the average number of seconds it takes to complete one iteration. It might seem surprising that the number of iterations required to obtain an approximately optimal solution decreases as we increase M . This phenomenon is driven by two factors. First, a larger value of M implies more optimal regular indicator matrices, hence more paths to optimality; Second, the initial indicator matrix we choose is more refined when M is larger. Due to these two countervailing forces, the total amount of time does not change dramatically as we increase M . These numerical results suggest that SIMS is quite scalable.

5.2.2. Performance Benchmark. Given that our problem is a convex optimization problem, it is useful to compare the speed of SIMS with a generic convex optimization solver. We choose three popular convex optimization packages: MOSEK, CVXOPT, and LOQO. MOSEK is a well-known commercial software for solving large-scale mathematical optimization problem, using the interior-point method. A recent survey compares MOSEK favorably to CPLEX, another leading commercial software for convex optimization (Ben-Tal and Nemirovski 2001). CVXOPT is a free python-based convex optimization software developed at UCLA, and LOQO is a commercial optimization software developed at Princeton University for smooth constrained optimization based on an infeasible, primal-dual, interior-point method.

We note that the interior-point method used by most commercial software requires the construction of the Hessian matrix during each iteration, which has a memory requirement in the order of $O((NM)^2)$.

Figure 3 Convergence of SIMS [Color figure can be viewed at wileyonlinelibrary.com]



In contrast, the memory requirement for SIMS is in the order of $O(NM)$ because all relevant variables during each iteration have the same dimension as the allocation matrix, which is $N \times M$. This implies commercial software such as MOSEK will have troubling fitting an exceedingly large problem into the memory. For this reason, we cap the problem size for MOSEK at $N = 100$ and $M = 5000$ so that it can finish within reasonable amount of time and the memory requirement.

We first compare the speed of MOSEK and SIMS by setting $N = 100$ and let M vary from 500 to 5000. The left panel of Figure 4 compares the time used by each software. Clearly, SIMS outperforms MOSEK when the scale of the problem is large. To compare the performance of SIMS with CVXOPT and LOQO, we further reduce the scale of NAS problems so that CVXOPT and LOQO can run properly. In particular, we set the number of agents to $N = 50$ and increase M from 50 to 100 at a step of 1. The right panel of Figure 4 compares the time used by each software. Clearly, the performance of SIMS is far superior to CVXOPT and LOQO.

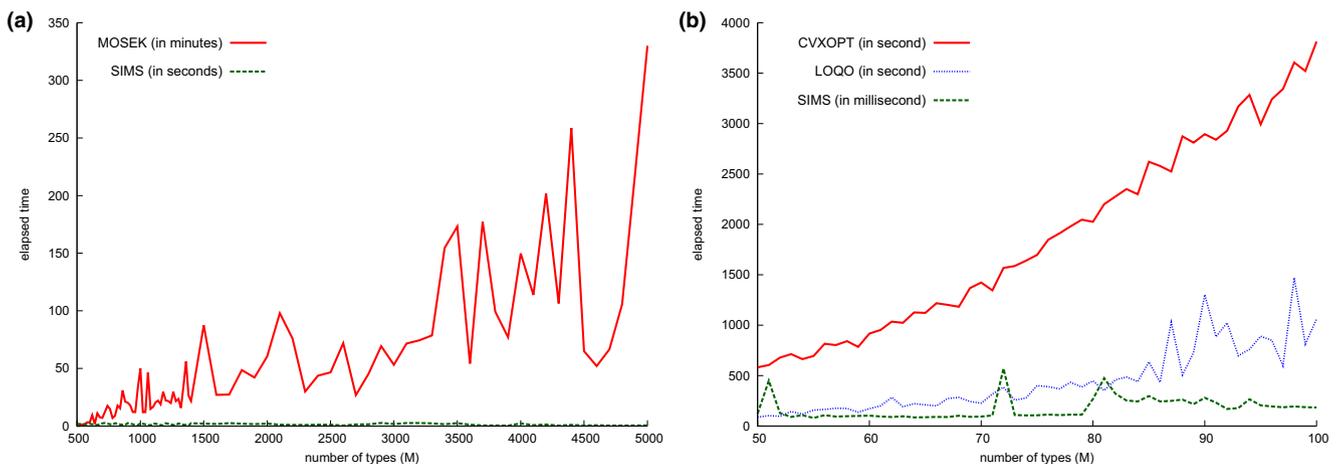
Based on these numerical experiments, we believe that SIMS has significant advantages in speed and memory requirement that make it particularly useful for solving extremely large-scale and/or time-critical problems: First, speed comparisons in Figure 4 suggest that for problems of large sizes, it would often take the standard algorithm hours to solve while it only takes SIMS a few seconds to solve. Second, because the memory requirement in SIMS is $O(NM)$ compared with $O((NM)^2)$ for most generic interior-point solvers, SIMS can solve much larger problems on commodity hardware, which by itself can justify the use of SIMS over generic convex optimization solvers.

5.2.3. Application to Display Advertising. To validate the applicability of the SIMS algorithm in real-world problems, we simulate a display advertising problem and use SIMS to solve them. Simulation methods have been used to test other algorithms for display advertising (Deza et al. 2015, Turner 2012). In display advertising, we reinterpret “agents” as ad campaigns to reflect the fact that each campaign has its own goals and preferences. Following Zhang et al. (2014) we assume that the each ad impression is characterized by K binary features (e.g., male/female, day/night, high income/low income, etc), resulting in 2^K total impression types. We also assume each campaign may target a small subset set of impression types, and different campaigns may use different features for targeting (e.g., one campaign may target gender while another may target income level). We discuss how we simulate supplies, targeting criteria for each campaign, and valuation coefficients below.

First, we let the number of features $K = 14$, resulting in 16,384 distinct impression types. Following Turner (2012), we use Pareto distribution to account for the fact that supplies are disproportionately large for some impression types. Specifically, for each impression feature, we draw two numbers, p_0 and p_1 , from a Pareto distribution with minimum 0, mean 1, and shape parameter 5. We then let $q_1 = \frac{p_1}{p_1 + p_0}$ be the probability of getting 1’s for this feature, and $q_0 = 1 - q_1$ for getting 0’s. We further assume that the probability of drawing an impression type with feature vector $\mathbf{f} = (f_1, f_2, \dots, f_K)$ is $P(\mathbf{f}) = q_{f_1}^1 q_{f_2}^2 \dots q_{f_K}^K$, where $q_{f_k}^k$ is the probability of drawing $f_k \in \{0, 1\}$ for feature k . We conduct random draws according to $P(\mathbf{f})$ to obtain the total supply of impression types.

Second, to simulate campaign targeting, we first draw a number k_1 from a Poisson distribution with

Figure 4 Comparison of SIMS with MOSEK (Left Panel), and with CVXOPT and LOQO (Right Panel)



Notes. The number of advertisers is fixed to 100 in the left panel and 50 in the right panel. [Color figure can be viewed at wileyonlinelibrary.com]

parameter 2 to be the number of targeted features. We then randomly choose k_1 features out of K as targeted features. For each targeted feature k , we let the targeting criterion be $f_k = 1$ with probability q_1^k , and $f_k = 0$ with probability q_0^k .

Finally, after simulating the targeting criteria for each campaign, we simulate the valuation coefficients for those targeted impression types. Assuming an exponential valuation function with parameters $\{v_i\}$ and $\{\alpha_{im}\}$, we randomly generate $\{\alpha_{im}\}$ for targeted impression types using a truncated normal distribution with mean 0.5, standard deviation 0.2, and min/max of 0.1 and 1, respectively. The coefficients are then scaled by a factor of 0.01 to make the optimization problem difficult enough¹⁷. Furthermore, to mimic the fact that campaigns have different budget levels, we simulate the coefficients $\{v_i\}$ by drawing from a truncated normal distribution with mean 0.5, standard deviation 0.2, and min/max of 0.1 and 1 respectively, and then scaling it by a factor of 10,000 to reduce floating-point numerical error although mathematically the scaled problem is essentially equivalent to the original one.

An upper bound of the NAS problem corresponding to this simulated real-world example is 5026061.52. SIMS solved this NAS problem in 446.945 seconds (roughly 7.5 minutes) with objective 5026061.519957. On the other hand, MOSEK failed to solve the problem within 100,000 iterations after 944904 seconds (roughly 11 days). The comparison suggests that the advantage of SIMS over MOSEK is enormous in more realistic scenarios.

6. Implications for Online Display Advertising

By solving the NAS problem, we can obtain several types of outputs: the price vector scaled by the Lagrange multiplier of the standardized problem, an optimal allocation, and a decomposition of the allocation matrix. The decomposition tells DSPs which audience categories and advertisers (ad campaigns) should be considered together. In the following, we focus on the implications of our two most important outputs: the price vector and the optimal allocation.

6.1. Implications of the Price Vector

We obtain a price vector as a by-product of the NAS problem, but because our model is rooted in economic theory, it has an intuitive economic interpretation and can be used in different ways. First, the price vector, as market clearing prices, can be used to determine a set of internal prices for DSP if DSPs were to charge these prices, advertisers should have no incentive to move away from the optimal allocation. Second, because the price vector has a shadow price

interpretation, DSPs can use these prices to decide whether it has too few or too many impressions for each audience category. For example, if the internal price for an audience category is higher than its market price, the DSP should consider buying more of such impressions.

6.2. Implication of the Optimal Allocation

A second, and probably more direct, application of our theory and algorithm is to guide the scheduling of display ads for DSPs. Our NAS problem can be part of the “optimize-and-dispatch” style ad scheduling system (Parkes and Sandholm 2005), where the first step is to solve an optimal NAS problem that produces an impression target for each campaign and audience category. Then an online dispatcher allocates incoming impressions one by one towards the impression targets. We briefly discuss below how the SIMS algorithm could be used for ad scheduling, including how to adapt to supply uncertainties.

Consider an environment where impressions arrive in a stochastic fashion over horizon $[0, T]$. We assume that there is an initial forecast about the total supply for each audience category and subsequent updated forecasts. Let $\omega^t = (\omega_1^t, \omega_2^t, \dots, \omega_m^t)$ be the forecasts for the total supply of all audience categories at time t . A general *optimize-and-dispatch* approach to ad scheduling can unfold like this:

1. (Initial optimization) Solve the NAS problem for the initial forecast ω^0 and obtain the initial optimal terminal allocation (i.e., the target number of impressions at T) \mathbf{x}^0 .
2. (Incremental optimization) At period t , if the forecast stays the same ($\omega^t = \omega^{t-1}$), we use the same terminal allocation $\mathbf{x}^t = \mathbf{x}^{t-1}$. Otherwise, we recompute \mathbf{x}^t as the solution to an NAS problem with the updated forecast ω^t .
3. (Dispatch) Allocate impressions upon arrival so that total allocated impressions are proportional to \mathbf{x}^t as much as possible.

Provided that the updated forecasts converge to the actual total supply as $t \rightarrow T$, the above optimize-and-dispatch procedure will approximate the optimal terminal allocation of the final NAS problem.

Now, what if the supply forecast changes? We believe that a SIMS-powered ad scheduling system can adapt to changing supplies fairly quickly. First, because an indicator matrix is optimal for a wide range of supply vectors, as long as the new forecast does not deviate much, we may not need a new indicator matrix. The only thing we need to do is to resolve a standardized NAS problem using the updated supply vector (steps 3–4 of Algorithm 2 in Appendix B.1), which can be done very efficiently. Even when the new forecast calls for a new indicator

matrix, we need not start from scratch because of the iterative nature of SIMS. We may simply iterate from the current indicator matrix still we reach a new optimal indicator matrix. Because the SIMS algorithm is shown to be very fast in our numerical experiments, such incremental iterations can be done fairly frequently (e.g., every 15 minutes).

To further improve the real-time performance of the SIMS algorithm, one may compute multiple probable supply scenarios ahead of time and store the solutions for later use. The SIMS algorithm permits us to store only the optimal indicator matrices and the associated price vectors, which can be combined with the supply forecast to quickly obtain the optimal allocation by solving a standardized NAS.

7. Conclusion

Motivated by real-world applications of online display advertising, we propose a unique class of allocation problem (NAS) where agents have concave value functions and different substitution preferences across numerous types of goods. Viewed as a transportation problem, our formulation permits greater flexibility in modeling agent preferences than existing transportation models because we allow multiple types of goods and agents to have heterogeneous rates of substitution for these goods.

Drawing upon the economic concept of Pareto optimality, we develop a theory and design an algorithm for solving NAS problem. The SIMS algorithm iterates through specially constructed indicator matrices each of which permits fast solution via a combination of decomposition and standardization techniques. Our simulation results show that SIMS runs up to three orders of magnitude faster than generic interior-point nonlinear solvers. Our theory has interesting connection with the martingale methodology used in asset pricing while our algorithm is connected to the Simplex algorithm for linear programming problems.

This research has its limitations that warrant further research. We have focused on non-physical goods and abstracted away transportation costs. It would be interesting to combine our problem with a transportation problem in a similar manner as Sharp et al. (1970). We have used exponential valuation functions for numerical experiments, it would be interesting to evaluate and compare the performance of SIMS under alternative valuation functions. So far, we have relied on numerical studies to establish the performance and scalability of SIMS. Future research could establish the complexity of SIMS. Once we have a regular indicator matrix, solving the corresponding single-good problem and verifying its optimality can be done quickly in polynomial time. We conjecture the iteration over regular indicator matrices to have

similar complexity as the iteration over vertices in the Simplex method, which is known to be exponential in the worst case but nevertheless takes polynomial time in practice (Spielman and Teng 2004).

This research can be extended in several ways. First, the current matrix search algorithm we use in SIMS is by no means the most efficient one and we believe it can be further improved with better heuristics. Because matrix regularity is an inherent property of any binary matrix, we hope future research on regularity can lead to powerful matrix search algorithms. Second, as we indicate in footnote 3, SIMS is applicable even if the objective functions are not concave. Its performance under non-concave objectives and comparison with generic nonlinear optimization software are promising directions for further study. Third, it would be interesting to both theoretically and numerically compare the SIMS algorithm with a recently proposed algorithm (Chubanov 2016) for separable convex optimization problems. Finally, it would also be appropriate to extend our problem to a stochastic setting. Extensive research has been done on decision under uncertainty, using stochastic programming (Sahinidis 2004, Shapiro et al. 2009) and robust programming (Bai et al. 1997, Mulvey et al. 1995). It would be interesting to explore the utility of our framework for dealing with resource allocation problems with heterogeneous preference for uncertainty.

Acknowledgments

We thank the department editor, Dr. Subodha Kumar, the anonymous senior editor, and the three anonymous reviewers for very constructive comments. We also thank seminar participants at UT Austin, Baidu Inc., Purdue University, Tsinghua University, Shanghai University of Finance and Economics, POMS Annual Conference, INFORMS Annual Meeting, Workshop on Information Systems and Economics, Workshop on Data Mining for Online Advertising, and Midwest Workshop on Control and Game theory for useful comments. Dr. De Liu received support for this research from National Science Foundation of China under Grant No. 71571044. Remaining errors are our sole responsibility.

Notes

¹For example, in Internet advertising, an Internet user's past behavior and geographic location can be tracked using browser cookies, allowing advertisers to draw inferences about a user's demographic background and interests. In mobile advertising, device and content characteristics, as well as geographic information may also be used to target and predict user interests.

²Practitioners seems to be ahead of the academics in terms of realizing the similarities between the two markets. For example, a co-founder of a digital ad trading company who spent 15 years in the financial industry commented that "We're talking about a market that shares a lot of the

same characteristics as financial markets” and they are “looking to apply investment banking tools and philosophies to online advertising.” For more details, please see the following Wall Street Journal article: <http://www.wsj.com/articles/SB10001424052702303949704579459103743176792>.

³Transportation problem is an important branch in the field of operations research, established several decades ago with pioneering works by Kantorovich (1960); Hitchcock (1941); Koopmans (1949); Dantzig (1951) and numerous subsequent contributions (see Ahuja et al. (1993) for a comprehensive overview).

⁴The entropy approach, for instance, imposes the same preference for diversity across all advertisers.

⁵For a comprehensive review of convex optimization, see Bazaraa et al. (2006); Boyd and Vandenberghe (2004).

⁶The number of impressions is typically extremely large in this industry, which makes the continuous relaxation of the decision variables less of a concern.

⁷The theory and the algorithm we will develop do not depend on the convexity assumption as long as certain technical requirements are met to ensure global optimality.

⁸Another way to generalize our formulation is to allow multiple portfolios per agent in a generalized valuation function. For example, for an agent who has 8 portfolios $S_{i1}, S_{i2}, \dots, S_{i8}$, we may define her valuation as

$$U_i(\mathbf{x}_i) = \sum_{l=1}^8 Q_{il} \left(\sum_{m \in S_{il}} \alpha_{im} x_{im} \right), \quad \forall i \in \mathcal{N}. \quad (4)$$

Treating each portfolio as a separate agents, we can clearly solve the generalized allocation problem in the same way as the original NAS problem. Because the marginal valuation for each portfolio decreases in the quantity allocated, an agent would prefer an allocation that spreads across multiple portfolios than those concentrate in one. In other words, such a generalized valuation function can capture the “variety-seeking” preferences.

⁹Our NAS problem can be converted into a separable convex optimization problem with general linear constraints, which is not NP-hard (Chubanov 2016). However, due to the high dimensionality of the solution space, general purpose convex optimization solvers, despite their theoretical efficiency, are not practical for solving large-scale NAS problems, as our numerical studies will show.

¹⁰We also use PO as a shorthand for *Pareto optimal*.

¹¹The exponential valuation function is commonly used in economics and finance (Malamud, et al. 2013, 2016) to capture an economic agent’s aversion to variation in consumption levels and the agent’s decreasing marginal utility from consumption.

¹²In the case of \mathbf{I}^1 , the 1’s in rows 1, 3, and 4 imply that $\frac{p_1}{p_2} = \frac{0.3}{0.16}$, $\frac{p_1}{p_3} = \frac{0.13}{0.4}$, $\frac{p_2}{p_4} = \frac{0.1}{0.3}$, which together yield a unique solution to \mathbf{p} (up to a scaling factor). In the case \mathbf{I}^2 , row 3 additionally implies $\frac{p_1}{p_4} = \frac{0.13}{0.08}$, which contradicts the existing conditions, thus a price vector does not exist.

¹³Note that $\tilde{Q}_i(\cdot)$ is well defined because for any agent i , when there are multiple m such that $I_{im} = 1$, we can use

any m to define $\tilde{Q}_i(\cdot)$ because the ratio α_{im}/p_m will be same for all different m by Equation (8).

¹⁴It should be noted that the existence of a solution to the system of linear Equations (12) is guaranteed by a technical result (Lemma 5) in Appendix A.

¹⁵A competitive equilibrium consists of a price vector and an allocation such that every agent prefers her current bundle to any other affordable bundle.

¹⁶Because $Q_i(\mathbf{x}_i) = v_i(1 - e^{-\sum_{m=1}^M \alpha_{im} x_{im}})$, one theoretical upper bound for the objective function is $\bar{Q} \equiv \sum_{i=1}^N v_i$.

¹⁷To see this, imagine the extreme case when the coefficients α_{ij} are extremely large. A trivial optimal allocation is to allocate enough supply to each campaign one by one so that its valuation approximates the upper bound (i.e., v_i). In general, the difficulty level of the problem increases with the scale of supply and the scale of the coefficients.

References

- Adler, M., P. B. Gibbons, Y. Matias. 2002. Scheduling space-sharing for Internet advertising. *J. Scheduling* 5(2): 103–119.
- Ahuja, R. K., T. L. Magnanti, J. B. Orlin. 1993. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, Englewood Cliffs, NJ.
- Babonneau, F., J.-P. Vial. 2009. ACCPM with a nonlinear constraint and an active set strategy to solve nonlinear multicommodity flow problems. *Mathemat. Program.* 120(1): 179–210.
- Bai, D., T. Carpenter, J. Mulvey. 1997. Making a case for robust optimization models. *Management Sci.* 43(7): 895–907.
- Bazaraa, M. S., H. D. Sherali, C. M. Shetty. 2006. *Nonlinear Programming: Theory and Algorithms*. John Wiley & Sons, Hoboken, NJ.
- Ben-Tal, A., A. Nemirovski. 2001. *Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications*, SIAM, Philadelphia, PA.
- Boyd, S., L. Vandenberghe. 2004. *Convex Optimization*, Cambridge University Press, Cambridge, UK.
- Chen, X., M. Sim, D. Simchi-Levi, P. Sun. 2007. Risk aversion in inventory management. *Oper. Res.* 55(5): 828–842.
- Chen, J., D. Liu, A. B. Whinston. 2009. Auctioning keywords in online search. *J. Market.* 73(4): 125–141.
- Chickering, D. M., D. Heckerman. 2000. Targeted advertising with inventory management. Proceedings of the 2nd ACM conference on Electronic commerce - EC '00, 145–149, ACM Press, New York, NY.
- Chubanov, S. 2016. A polynomial-time descent method for separable convex optimization problems with linear constraints, *SIAM J. Optimiz.* 26(1) 856–889.
- Dantzig, G. B. 1951. Application of the simplex method to a transportation problem. T. C. Koopmans, ed. *Activity Analysis and Production and Allocation*. John Wiley and Sons, New York, 359–373.
- Duffie D. 2001. *Dynamic Asset Pricing Theory*, Princeton University Press, Princeton, NJ.
- Deane, J., A. Agarwal. 2012. Scheduling online advertisements to maximize revenue under variable display frequency. *Omega* 40(5): 562–570.
- Deza, A., K. Huang, M. R. Metel. 2015. Chance constrained optimization for targeted Internet advertising. *Omega* 53: 90–96.
- EMarketer. 2016. US digital display ad spending to surpass search ad spending in 2016. Available at <https://www.emarketer.com/Article/US-Digital-Display-Ad-Spending-Surpass-Search-Ad-Spending-2016/1013442> (accessed date December 22, 2016).

- Gul, F., E. Stacchetti. 1999. Walrasian equilibrium with gross substitutes. *J. Econ. Theory* 87(1): 95–124.
- Harrison, J. M., D. M. Kreps. 1979. Martingale and arbitrage in multiperiod securities markets. *J. Econ. Theory* 20(3): 381–408.
- Hitchcock, F. L. 1941. The distribution of a product from several sources to numerous facilities. *J. Mathemat. Phys.* 20(1–4): 224–230.
- Ho, C.-J., J. Vaughan. 2012. Online task assignment in crowd-sourcing markets. Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, 45–51, July 22–26, Toronto, Ontario, Canada.
- Hochbaum, D. S. 2007. Complexity and algorithms for nonlinear optimization problems. *Ann. Oper. Res.* 153(1): 257–296.
- Hochbaum, D. S., J. G. Shanthikumar. 1990. Convex separable optimization is not much harder than linear optimization. *J. Assoc. Comput. Mach.* 37(4): 843–862.
- Kantorovich, L. V. 1960. Mathematical methods in the organization and planning of production. *Management Sci.* 6(4): 366–422.
- Katoh, N., T. Ibaraki. 1998. Resource allocation problems. D. Du, P. M. Pardalos, eds. *Handbook of Combinatorial Optimization*, Kluwer Academic Publishers, Norwell, MA, 159–260.
- Koopmans, T. C. 1949. Optimum utilization of the transportation system. *Econometrica* 17: 136–146.
- Kumar, S., V. S. Jacob, C. Sriskandarajah. 2006. Scheduling advertisements on a web page to maximize revenue. *Eur. J. Oper. Res.* 173(3): 1067–1089.
- Kumar, S., M. Dawande, V. Mookerjee. 2007. Optimal scheduling and placement of Internet banner advertisements. *IEEE Trans. Knowl. Data Eng.* 19(11): 1571–1584.
- Lahaie, S., D. C. Parkes, D. M. Pennock. 2008. An expressive auction design for online display advertising. Proceedings of the Twenty-third AAAI Conference on Artificial Intelligence, 108–113, July 13–17, 2008, Chicago, Illinois.
- Langheinrich, M., A. Nakamura, N. Abe, T. Kamba, Y. Koseki. 1999. Unintrusive customization techniques for Web advertising. *Comput. Network* 31(11–16): 1259–1272.
- Liu, D., S. Viswanathan. 2014. Information asymmetry and hybrid advertising. *J. Market. Res.* 51(5): 609–624.
- Malamud, S., H. Rui, A. Whinston. 2013. Optimal incentives and securitization of defaultable assets. *J. Financ. Econ.* 107(1): 111–135.
- Malamud, S., H. Rui, A. Whinston. 2016. Optimal reinsurance with multiple tranches. *J. Mathemat. Econ.* 65(August): 71–82.
- Markowitz, H. M. 1959. *Portfolio Selection: Efficient Diversification of Investments*. John Wiley & Sons, New York, NY.
- McAfee, P. 2011. The design of advertising exchanges. *Rev. Ind. Organ.* 39(3): 169–185.
- McAfee, P., K. Papineni, S. Vassilvitskii. 2010. Maximally representative allocation for guaranteed delivery advertising campaigns. *Rev. Econ. Des.* 17(2): 83–94.
- MOSEK. 2015. *The MOSEK Optimization Tools Manual*, MOSEK ApS, Denmark.
- Mulvey, J. M., R. J. Vanderbei, S. A. Zenios. 1995. Robust optimization of large-scale systems. *Oper. Res.* 43(2): 264–281.
- Muthukrishnan, S. 2009. Ad exchanges: Research issues. S. Leonard, ed. *Lecture Notes in Computer Science* (Volume 5929). Springer, Berlin, Heidelberg, 1–12.
- Nakamura, A., N. Abe. 2005. Improvements to the linear programming based scheduling of Web advertisements. *Electron. Commer. Res.* 5(1): 75–98.
- Ouorou, A. 2007. Implementing a proximal algorithm for some nonlinear multicommodity flow problems. *Networks* 49(1): 18–27.
- Ouorou, A., P. Mahey, J.-P. Vial. 2000. A survey of algorithms for convex multicommodity flow problems. *Management Sci.* 46(1): 126–147.
- Parkes, D. C., T. Sandholm. 2005. Optimize-and-dispatch architecture for expressive ad auctions. EC '05: Proceedings of the 6th ACM Conference on Electronic Commerce, Vancouver, Canada, June 5–8, 2005. Association for Computing Machinery. New York, NY.
- Patriksson, M. 2008. A survey on the continuous nonlinear resource allocation problem. *Eur. J. Oper. Res.* 185(1): 1–46.
- Ross, S. A. 1978. A simple approach to the valuation of risky streams. *J. Bus.* 51(3): 453–475.
- Sahinidis, N. V. 2004. Optimization under uncertainty: State-of-the-art and opportunities. *Comput. Chem. Eng.* 28(6–7): 971–983.
- Shapiro, A., D. Dentcheva, A. Ruszczyński. 2009. *Lectures on Stochastic Programming: Modeling and Theory*. SIAM-Society for Industrial and Applied Mathematics, Philadelphia, PA.
- Sharp, J. F., J. C. Snyder, J. H. Greene. 1970. A decomposition algorithm for solving the multifacility production-transportation problem with nonlinear production costs. *Econometrica* 38(3): 490–506.
- Shetty, C. M. 1959. A solution to the transportation problem with nonlinear costs. *Oper. Res.* 7(5): 571–580.
- Spielman, D., S. Teng. 2004. Smoothed analysis of algorithms: Why the Simplex algorithm usually takes polynomial time. *J. Assoc. Comput. Mach.* 51(3): 385–463.
- Tomlin, J. A. 2000. An entropy approach to unintrusive targeted advertising on the Web. *Comput. Netw.* 33(1–6): 767–774.
- Turner, J. 2012. The planning of guaranteed targeted display advertising. *Oper. Res.* 60(1): 18–33.
- Vanderbei, R. J. 1997. *LOQO User's Manual*.
- Vidakovic, R. 2013. 5 ways real-time bidding differs from direct buys. Available at <http://marketingland.com/5-ways-real-time-bidding-differs-from-direct-buys-37384> (accessed date December 22, 2016).
- Wilbur, K. C., L. Xu, D. Kempe. 2013. Correcting audience externalities in television advertising. *Market. Sci.* 32(6): 892–912.
- Ye, H., D. Yao. 2010. Utility-maximizing resource control: Diffusion limit and asymptotic optimality for a two-bottleneck model. *Oper. Res.* 58(3): 613–623.
- Zhang, W., S. Yuan, J. Wang, X. Shen. 2014. Real-time bidding benchmarking with iPinYou dataset. Working paper. arXiv:1407.7073

Supporting Information

Additional supporting information may be found online in the supporting information tab for this article:

Appendix A: Proofs.

Appendix B: Additional Details of SIMS.