



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Value Loss in Allocation Systems with Provider Guarantees

Yonatan Gur, Dan Iancu, Xavier Warnes

To cite this article:

Yonatan Gur, Dan Iancu, Xavier Warnes (2021) Value Loss in Allocation Systems with Provider Guarantees. Management Science 67(6):3757–3784. <https://doi.org/10.1287/mnsc.2020.3656>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2020, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Value Loss in Allocation Systems with Provider Guarantees

Yonatan Gur,^a Dan Iancu,^{a,b} Xavier Warnes^a

^a Graduate School of Business, Stanford University, Stanford, California 94305; ^b INSEAD, Fontainebleau 77300, France

Contact: ygur@stanford.edu,  <https://orcid.org/0000-0003-0764-3570> (YG); daniancu@stanford.edu,

 <https://orcid.org/0000-0001-9890-0928> (DI); xwarnes@stanford.edu,  <https://orcid.org/0000-0001-5436-9362> (XW)

Received: March 21, 2019

Revised: December 17, 2019; March 24, 2020

Accepted: April 9, 2020

Published Online in Articles in Advance:
October 13, 2020

<https://doi.org/10.1287/mnsc.2020.3656>

Copyright: © 2020 INFORMS

Abstract. Many operational settings share the following three features: (i) a centralized planning system allocates tasks to workers or service providers, (ii) the providers generate value by completing the tasks, and (iii) the completion of tasks influences the providers' welfare. In such cases, the planning system's allocations often entail trade-offs between the service providers' welfare and the total value that is generated (or that accrues to the system itself), and concern arises that allocations that are good under one metric may perform poorly under the other. We propose a broad framework for quantifying the magnitude of value losses when allocations are restricted to satisfy certain desirable guarantees to the service providers. We consider a general class of guarantees that includes many considerations of practical interest arising (e.g., in the design of sustainable two-sided markets) in workforce welfare and compensation, or in sourcing and payments in supply chains, among other application domains. We derive tight bounds on the relative value loss and show that this loss is limited for any restriction included in our general class. Our analysis shows that when many providers are present, the largest losses are driven by fairness considerations, whereas when few providers are present, they are driven by the heterogeneity in the providers' effectiveness to generate value; when providers are perfectly homogenous, the losses never exceed 50%. We study additional loss drivers and find that less variability in the value of jobs and a more balanced supply-demand ratio may lead to larger losses. Lastly, we demonstrate numerically using both real-world and synthetic data that the loss can be small in several cases of practical interest.

History: Accepted by Chung Piaw Teo, optimization.

Supplemental Material: Data are available at <https://doi.org/10.1287/mnsc.2020.3656>.

Keywords: allocation systems • efficiency loss • worst-case analysis • worker guarantees • fairness

1. Introduction

In many operational settings, a centralized planning system (*system* henceforth) decides how to allocate a pool of resources or tasks to an existing set of workers or service providers. Such allocations routinely determine the total value created; but oftentimes, they also influence how this value is shared between the system and the providers. When only some allocations meet particular attributes that the providers find appealing, this could significantly impact both the providers' welfare as well as the system in the long run (e.g., because of provider retention).

This gives rise to a potentially difficult question: how should the system trade off the appealing features of such allocations against the potential value loss associated with them? Critically, how much of the total value created or of the system's share of that value might be lost by ensuring that allocations to providers meet such desirable attributes?

To understand this fundamental issue in a more concrete setting, consider online service platforms such as Upwork, Grubhub, Uber, or Lyft. These platforms match customer service requests for labor,

food, or rides, with dedicated service providers (freelance workers or drivers), and typically rely on revenue-sharing agreements to split the revenue collected between the providers and the platform. The allocation of service requests thus critically influences the value generated upon each service completion as well as the portions of the total value that are retained by different providers. This value may consist of both monetary components (e.g., the revenue collected, or the profits when accounting for costs) as well as nonmonetary ones (e.g., the satisfaction of providers and the service quality experienced by customers). *Ceteris paribus*, providers who are assigned fewer or "worse" requests could end up with lower welfare and as a consequence, could potentially leave the platforms for better prospects. Such retention issues have been well documented (e.g., CNBC 2017), and platforms have begun to set up a variety of mitigating measures that range from guaranteed income levels for providers (see, e.g., Lyft Inc. 2018, Uber Technologies Inc. 2018b, and GrubHub 2019) to designing loyalty and bonus payments tied to completing multiple service requests (see, e.g., *Financial Times* 2014, Uber Technologies Inc. 2018a, and DoorDash 2019).

How much are platforms sacrificing in terms of revenues, profits, or customer experience when their allocations are designed to carry such guarantees for the service providers?

To demonstrate the trade-off in a different setting, consider a traditional brick-and-mortar business that designs work schedules for its sales associates. Here too the allocation can critically drive the value generated: assigning a top-performing sales associate to work on busier days may increase sales and revenue for the store, and may also increase customer satisfaction. But, such allocations also have a direct impact on the employees, in both monetary terms (e.g., because of commissions and bonuses tied to completed sales) (Berger 1972) as well as nonmonetary terms (e.g., because of job satisfaction, work-life balance, and worker health) (Bacharach et al. 1991, Sparks et al. 2011). Work schedules and job assignments thus routinely follow certain patterns intended to maintain a fair and balanced workload for the employees. But, do these entail a significant revenue or profit loss for the employer or goodwill loss for the customers?

These examples highlight several settings in which only some of the system's allocations meet certain desirable attributes for the service providers. We henceforth refer to these as *provider guarantees*. Considering only allocations that ensure provider guarantees could generate a loss in the total value or in the system's share of the value that could be achieved without such restrictions. Although such losses could in theory be mitigated in some cases by designing suitable monetary transfers,¹ the resulting mechanisms are rarely implementable in practice because of numerous legal, ethical, and computational challenges.² We thus focus on understanding the value loss associated with provider guarantees in settings where monetary transfers are not possible; we seek to quantify the magnitude of this value loss, its key drivers, and the structure of the guarantees that are most likely to cause large losses.

1.1. Main Contribution

On the modeling front, we develop a broad framework that allows quantifying the value that may be lost because of imposing provider guarantees in various settings. We consider a centralized planning system that allocates a discrete set of jobs (or resources) to a set of providers who are endowed with heterogeneous effectiveness to generate value by completing the jobs. We capture an allocation design that institutes desirable provider guarantees by imposing constraints (also referred to as *restrictions*) on the system's feasible allocations, and we only require the set of constrained allocations to satisfy a mild and natural monotonicity condition. This allows us to capture a variety of practical considerations arising in the sustainable

design of two-sided markets, in workforce welfare and compensation mechanisms, as well as in sourcing and payments in supply chains, among other application domains.

We define the relative value loss associated with instituting certain provider guarantees as the fraction of the maximal value achievable by unrestricted allocations that is lost when imposing the constraints associated with satisfying these guarantees. We derive tight bounds on the relative value loss that hold for *any* restrictions in a general class of provider guarantees. We establish these by solving a fractional linear relaxation of the problem of maximizing the relative value loss, and by producing instances that match the maximal value of this relaxation.

Our bounds only depend on the heterogeneity in the providers' effectiveness to generate value, and on the number of providers. We find that although unbounded heterogeneity may lead to unbounded relative loss, when heterogeneity is bounded, the loss must be bounded as well under any of the guarantees that we consider. In particular, the relative loss never exceeds 50% when providers are homogenous in their effectiveness to generate value, irrespective of how many providers exist. This contrasts several findings in the resource allocation literature that exhibit unbounded losses, particularly in regimes with a large number of providers, and highlights the importance of explicitly capturing certain aspects of the job allocation problem that are included in our model (see Section 1.2 for a more detailed discussion). Qualitatively, our findings suggest that allocation systems concerned with large losses because of instituting provider guarantees may consider reducing the heterogeneity in the providers' (endowed) effectiveness to generate value, as this can substantially reduce the magnitude of such losses in the worst case.

Our analysis also highlights the most prominent settings and factors leading to value loss. We show that when few providers exist, the key loss driver is the *heterogeneity* in the providers' effectiveness to generate value. The largest losses emerge when the allocation system institutes guarantees that target precisely the providers with lower effectiveness to generate value; in practice, this could occur, for example, when guaranteeing a certain workload level to new, inexperienced providers.

In contrast, we show that when many providers with bounded heterogeneity exist, the primary loss driver is *lost demand* (i.e., unallocated jobs): the largest losses occur when the system is unable to allocate many of the jobs because of a combination of physical constraints (preventing jobs from being completed by the same provider) and fairness constraints related to the guarantees instituted for the providers. More precisely, we find that the worst-case losses emerge

from max-min fairness considerations. Such fairness considerations are a strict subset of the guarantees that we consider; thus, our result provides additional context for a recent literature stream focused on quantifying fairness-efficiency trade-offs in operations (see Section 1.2 for a more detailed discussion), by showing that these may become particularly pertinent in large-scale allocation systems with ample supply, ample demand, and physical constraints that prevent certain jobs from being assigned together.

We also study the impact of several other important drivers of the value loss. Building on the observation that the structure of the set of feasible allocations can be critical, we characterize several cases of practical interest where the loss is guaranteed to vanish. In addition, we show that the integrality of allocations is critical: allowing for fractional allocations would eliminate the loss for a broad class of provider guarantees. We further show that the symmetry of the set of feasible allocations plays a prominent role: when providers have different sets of jobs they can perform, the relative value loss may asymptotically approach 100% as the number of providers grows large. This is consistent with our finding concerning the impact of heterogeneity in the providers' effectiveness to generate value, and it reinforces the idea that when providers are sufficiently "different" (either in their effectiveness to generate value from jobs or in their capability to execute jobs in the first place), this can critically drive the losses, making them unbounded in extreme cases.³ Finally, we demonstrate that the variability in the value of jobs and the imbalance in supply (number of providers) and demand (number of jobs) significantly impact value losses, with a higher variability in values or a more imbalanced supply-demand leading to reduced value losses.

Using both synthetic data and real-world data consisting of taxi trips in New York City, we study numerically the relative value losses associated with implementing provider guarantees. As this setting corresponds roughly to viewing providers (i.e., taxi drivers) as approximately homogeneous in their effectiveness to generate value from jobs, we focus on guarantees corresponding to fairness considerations. The numerical results confirm the robustness of our theoretical findings concerning the impact of the variation in the intrinsic job value and the supply-demand imbalance on the relative value loss. Moreover, in the instances generated from the real-world data, we document relative value losses that never exceed 4%. This suggests that losses associated with implementing provider guarantees in particular cases of practical interest may actually be even significantly lower than our theoretical worst-case bounds—a finding that we believe motivates additional research focused on more specific settings or on finding policies that can achieve such guarantees.

1.2. Related Literature

1.2.1. Relative Efficiency Losses Because of Fairness. Our work is related to a stream of literature that studies efficiency losses emerging when outcomes are constrained to satisfy certain fairness considerations. Bertsimas et al. (2011, 2012) consider continuous resource allocation problems where a centralized decision maker balances efficiency (i.e., social welfare) with fairness and equity considerations. They define the *price of fairness* as the relative loss in efficiency under such fairness considerations and derive theoretical bounds on this measure that depend on the number of agents and on the fairness criterion. As the number of agents grows arbitrarily large, the worst-case losses always approach 100%, regardless of the fairness criterion imposed, and the instances that get close to this loss involve asymmetries between the different agents. Our study retains the same definition of relative value loss but differs in several modeling primitives and results. Specifically, we focus on *discrete* allocation problems where guarantees (such as fairness considerations) are modeled through constraints on the feasible allocations rather than on the possible utility outcomes; we consider a broader set of restrictions that includes the fairness criteria in Bertsimas et al. (2011, 2012) as special cases; and we explicitly model and study a specific form of heterogeneity, in the providers' effectiveness to generate value from the jobs they are assigned. In our setting, unbounded heterogeneity is needed to obtain losses of 100%, and this can occur with both a small and a large number of agents/providers, unlike in Bertsimas et al. (2011, 2012). More importantly, we find that when heterogeneity is bounded, the worst-case losses are always bounded irrespective of the number of providers and of the guarantees used, and never exceed 50% when providers are homogenous. We elaborate more on the root causes for this discrepancy in Section 4. Our finding that worst-case restrictions actually correspond to max-min fairness considerations when the number of agents/providers is sufficiently large also provides more context for the large stream of literature focused on quantifying fairness and efficiency trade-offs in various large-scale operational settings (e.g., Bertsimas et al. 2013, Iancu and Trichakis 2014, McCoy and Lee 2014, Qi 2017).

1.2.2. Efficiency Losses in Equilibrium. Our paper is also related to a rich literature studying the *Price of Anarchy*—a measure introduced by Papadimitriou (2001) and Roughgarden and Tardos (2002) that quantifies the efficiency loss of Nash equilibrium outcomes relative to an optimal centralized solution. It is known that the Price of Anarchy can be bounded in certain settings (see, e.g., Roughgarden 2003, Correa et al. 2004, Johari and Tsitsiklis 2004, Perakis and Roels 2007),

but it can also be arbitrarily large (see, e.g., Awerbuch et al. 2006, Chawla and Roughgarden 2008, Koutsoupias and Papadimitriou 2009). We study efficiency losses generated by a centralized planner restricting the outcomes on purpose, rather than losses because of strategic behavior of selfish agents.

1.2.3. Allocation of Indivisible Jobs. Several approximation algorithms have been proposed to obtain envy-free and max-min fair allocations for *indivisible* goods (see, e.g., Lipton et al. 2004, Golovin 2005, and Asadpour and Saberi 2010). This line of work is aimed at determining the allocations themselves, whereas our paper is focused on quantifying the losses associated with such allocations, and understanding their key drivers.

1.2.4. Efficiency of Contracts. Our work is also related to a body of literature studying the efficiency losses that may arise in various principal-agent interactions, such as between a firm's shareholders, debtholders, and managers (see, e.g., Jensen and Meckling 1976), between firms and their sales associates (see, e.g., Farley 1964), between buyers and their suppliers (see, e.g., Cachon and Lariviere 2005), etc. Several papers in this literature also seek to quantify the associated efficiency losses (see, e.g., Besbes et al. 2018 for more discussion and additional references). Our work is not concerned with agency considerations but instead, focuses on quantifying the relative efficiency loss associated with restricting allocations to satisfy certain attributes that may be desirable to providers.

2. Problem Formulation

For the sake of clarity, we first provide a basic description of our setup and then discuss some concrete examples in Section 2.2. A discussion of the modeling assumptions is deferred to Section 2.3.

Consider a centralized planning system that allocates a given set of jobs D to a set of n service providers denoted by $N = \{1, \dots, n\}$. Each job possesses a certain *intrinsic value*, which we capture through a function $v : D \rightarrow \mathbb{R}$, so that $v(d)$ denotes the intrinsic value for job $d \in D$. For any subset of jobs $S \subseteq D$, we denote by $v(S) \stackrel{\text{def}}{=} \sum_{d \in S} v(d)$ the total value of all the jobs in S . Not all allocations of jobs to providers are possible, and we let \mathcal{F} denote the set of *feasible allocations* of jobs in D . If $A = (A_1, \dots, A_n) \in \mathcal{F}$ denotes a feasible allocation, then A_i denotes the jobs allocated to provider $i \in N$, and $A_{-i} \stackrel{\text{def}}{=} (A_1, \dots, A_{i-1}, A_{i+1}, \dots, A_n)$ denotes the allocation to all other providers.

When provider i is assigned a set of jobs $A_i \subseteq D$, the value that is generated is $\gamma_i v(A_i)$. The parameter γ_i is predetermined and fixed, and belongs to an interval $[\gamma_{\min}, \gamma_{\max}]$, where $0 < \gamma_{\min} \leq \gamma_{\max}$.⁴ We let $\gamma \in [\gamma_{\min}, \gamma_{\max}]^n$ denote the heterogeneity profile of

providers, and we denote the degree of heterogeneity by

$$\delta := \frac{\gamma_{\max} - \gamma_{\min}}{\gamma_{\max}}.$$

2.1. System Considerations and Value Loss

In deciding the allocations, the system seeks to generate as much value as possible. Given all the feasible allocations \mathcal{F} , the allocation that would maximize the total value generated would be given by the optimal solution to the problem

$$\max_{A \in \mathcal{F}} \sum_{i=1}^n \gamma_i v(A_i).$$

In order to guarantee certain conditions to its service providers, the system would also be interested in restricting attention to a subset of *allocations with guarantees* $\mathcal{F}_G \subseteq \mathcal{F}$. This could reduce the total value generated, and we define the *value loss under provider guarantees* $L_\gamma(\mathcal{F}, \mathcal{F}_G)$ as the relative loss in total value when the system considers only allocations from \mathcal{F}_G :

$$L_\gamma(\mathcal{F}, \mathcal{F}_G) = \frac{\max_{A \in \mathcal{F}} \sum_{i=1}^n \gamma_i v(A_i) - \max_{B \in \mathcal{F}_G} \sum_{i=1}^n \gamma_i v(B_i)}{\max_{A \in \mathcal{F}} \sum_{i=1}^n \gamma_i v(A_i)}. \quad (1)$$

Throughout the paper, we restrict attention to cases where \mathcal{F}_G is nonempty. Our goal is to understand the magnitude and the key drivers of this value loss. The general model setup allows capturing many practical settings of interest, as we discuss next.

2.2. Examples

We next illustrate the interpretation of various model components using a series of practical examples. In Table 1, we summarize the interpretations of these components.

2.2.1. Service Platforms. Service platforms such as Upwork, Grubhub, Uber, and Lyft can be thought of as systems that allocate demand for services (i.e., labor, food delivery, trips) to providers (i.e., freelance workers, drivers). The jobs are indivisible, and each job has a certain intrinsic value $v(d)$, with several possible interpretations. For instance, $v(d)$ may correspond to the revenue from completing job d (given by the amount paid by the customer), in which case the heterogeneity parameter γ_i can capture different notions of generated value, including the following examples.

i. *Total revenue.* Taking $\gamma_i = 1$ for all i , the objective $\sum_{i=1}^n \gamma_i v(A_i)$ captures the total revenue from the completed jobs, a measure of economic efficiency in the absence of costs.

ii. *Revenue sharing.* It is customary for service platforms such as Upwork, Grubhub, Uber, or Lyft to retain a fraction of the revenue generated from the provided services. In this case, if γ_i represents the share of the revenue accruing to the platform when dealing with provider i , then $\sum_{i=1}^n \gamma_i v(A_i)$ would represent the total revenue collected by the platform.

iii. *Costs/profits.* The spatial and temporal lengths of a job are typically key drivers for both the revenues as well as the costs from completing the job. Thus, $\gamma_i v(d)$ could capture the gross profit when driver i completes the job, equal to the revenue $v(d)$ net of provision costs $(1 - \gamma_i)v(d)$. This allows modeling heterogeneity in the transportation costs, for example, because of the different fuel economy of cars. The total value generated $\sum_{i=1}^n \gamma_i v(A_i)$ would denote the total net profit from jobs, a measure of economic efficiency in the presence of costs.

In addition to these measures, if $v(d)$ denotes the spatial or temporal length of a job d , then $\gamma_i v(d)$ could also capture the quality of service experienced by the customer(s) when provider i completes the job. Under this interpretation, γ_i may capture idiosyncratic differences because of, for example, car cleanliness or driver friendliness. The total generated value would then correspond to the total quality of service experienced by customers from the allocations.

The set of feasible allocations \mathcal{F} captures constraints on allocating jobs to providers: for instance, that each trip can be allocated to at most one provider, and that two trips that overlap in time cannot be allocated together to the same driver. The allocations with guarantees $\mathcal{F}_G \subseteq \mathcal{F}$ can capture, for example, a platform's commitments for minimal providers' income (see, e.g., Uber Technologies Inc. 2018b). In this setting, the value loss could thus be driven by two key factors: the providers' heterogeneous effectiveness to generate (monetary or nonmonetary) value from the jobs they are allocated and the potential demand loss because of the inability to allocate all jobs.

2.2.2. Workforce Management. A closely related example occurs when managing a workforce such as a team of sales associates. In this case, the system could be a regional or store manager who designs schedules for n sales associates. These schedules could be temporal (e.g., which hours or days to work), spatial (e.g., which floors or departments to cover), or could consist of the assignments of particular clients. With $v(d)$ denoting the (expected) revenue from a particular sales opportunity d , $\gamma_i v(d)$ can capture the revenue generated when this opportunity is assigned to associate i (with γ_i measuring the associate's effectiveness/performance), or it could capture the fraction of the revenue accruing to the firm (with $1 - \gamma_i$ denoting associate i 's commission). As in the service platform

example, the total value could also capture the quality of service experienced by clients or the gross profit when the associates incur variable costs; the set of feasible allocations \mathcal{F} could capture constraints on the schedules or assignments; and the subset of allocations $\mathcal{F}_G \subseteq \mathcal{F}$ could capture guarantees in terms of income, bonuses, or even spare time (all of which are known to be relevant to effective workforce management) (see, e.g., Tremblay et al. 2000, Cohen-Charash and Spector 2001).

2.2.3. Sourcing from a Heterogeneous Supply Base. A different example arises when the system represents a firm that decides how to allocate prescheduled indivisible orders for inputs among its n suppliers. Here, $v(d)$ can capture the volume of a particular order d , and the coefficients γ_i may capture supply yields (when suppliers are heterogeneous in their reliability, effectiveness, or quality) or gross margins for the firm (when different prices are paid to different suppliers). The set of allocations $\mathcal{F}_G \subseteq \mathcal{F}$ may capture guarantees for income or workload that could arise from a variety of considerations, such as a long-term sourcing strategy that requires keeping multiple qualified suppliers, implementing dual-sourcing policies (Yu et al. 2009, Yang et al. 2012), or particular social or environmental responsibility commitments (Patagonia 2018, Starbucks Corporation 2018).

2.3. Assumptions

The allocation problem we described so far is very general but is also intractable in the absence of additional structure. To that end, we next introduce some mild assumptions that still permit a lot of generality in the allowable primitives of our model and yet, render tractability in settings of practical interest. The first assumption concerns jobs and feasible allocations, and the second concerns the provider guarantees that can be under consideration. Throughout, we use $\mathcal{P}(D)$ to denote the collection of all subsets of D .

Assumption 1. *The set of feasible allocations \mathcal{F} satisfies the following properties:*

i. (*Indivisibility*) *In any feasible allocation, each job is assigned to at most one provider: that is,*

$$\mathcal{F} \subseteq \{A = (A_1, \dots, A_n) \in \mathcal{P}(D)^n \mid A_i \cap A_j = \emptyset \text{ for all } i, j \in N, i \neq j\}.$$

ii. (*Symmetry*) *If A is a feasible allocation, then any permutation of A is a feasible allocation.*

iii. (*Monotonicity*) *If (A_i, A_{-i}) is feasible, then (B, A_{-i}) is feasible, for any $B \subseteq A_i$ and any $i \in N$.*

iv. (*Provider Independence*) *If $(A_i, A_{-i}), (B_i, B_{-i}) \in \mathcal{F}$ are such that $A_i \cap B_j = \emptyset$ for all $j \neq i$, then $(B_1, \dots, B_{i-1}, A_i, B_{i+1}, \dots, B_n) \in \mathcal{F}$.*

Table 1. Interpretation of Parameters

$v(d)$	γ_i	$\sum_{i=1}^n \gamma_i v(A_i)$	\mathcal{F}	\mathcal{F}_G
Service platforms				
Revenue from job d	Revenue generated by provider i	Total revenue	Constraints in allocating jobs (e.g., each trip can be allocated to at most one provider)	Platform's commitments to the providers (e.g., minimum income)
Revenue from job d	Platform's share of revenue generated by provider i	Total revenue accrued by platform		
Revenue from job d	Gross profit generated by provider i	Total net profit		
Length of job d	Quality of service offered by provider i	Total quality of service experienced by customers		
Workforce management				
Expected revenue from job d	Associate i 's effectiveness/performance	Total expected revenue	Constraints on the schedules or assignments	Guarantees in terms of income, bonuses, or spare time
Sourcing from heterogeneous suppliers				
Volume of order d	Supply yields or gross margins for the firm	Total value generated by the suppliers	Intrinsic constraints in the allocation of orders to suppliers	Dual-sourcing policies, social or environmental responsibility commitments

Notes. For each example in Section 2.2, we present a summary of the interpretation of each parameter.

Part (i) of Assumption 1 requires that jobs are *indivisible*, so that feasible allocations can assign each job to at most one provider. This is reasonable in various settings, including (a) service platforms such as Upwork, Grubhub, Uber, or Lyft, where unique jobs must be allocated to service providers operating independently; (b) sales settings where a single lead cannot be divided across multiple associates; and (c) sourcing settings where a single unit cannot be divided among suppliers.

Part (ii) of the assumption implicitly requires providers to be homogeneous in their ability to perform jobs: if a set of jobs can be fulfilled by one provider, it can be fulfilled by any other provider as well. This is reasonable when the jobs do not require essential skills or technology that is available only to a subset of suppliers or service providers. Thus, it would not hold in settings in which providers' specialization plays a prominent role in their ability to complete jobs. Nevertheless, it is important to note that this requirement only pertains to *feasibility* (i.e., it does not imply that different providers generate the same value when completing a particular job). For instance, different Upwork providers may all be able to complete a particular task, but the value that is generated (e.g., the profit or the quality of service experienced by the customer) may differ, which could be captured by the coefficients $\{\gamma_i\}_{i=1}^n$.

Part (iii) states that it is always possible to allocate fewer jobs. This always holds when jobs can be carried

out independently from one another. In addition, the requirement allows certain dependencies between jobs: for instance, if a set of jobs must be completed together (i.e., by a single provider), these could be grouped into a single aggregate job that should be allocated as an indivisible object in our setup.

Finally, part (iv) states that feasible allocations can be obtained by concatenating feasible allocations for subsets of providers, as long as no job is assigned more than once. This is essentially a requirement of independence on the providers: as long as jobs are not duplicated, whether a provider can fulfill a set of jobs is independent of what jobs the other providers are fulfilling. This is reasonable in many settings where having one provider complete certain jobs carries essentially no externalities on other providers.

Our last assumption provides a connection between the set of allocations with guarantees \mathcal{F}_G and the value of jobs v .

Assumption 2 (Pareto-Efficient Provider Guarantees). *If $B \in \mathcal{F}_G$ and $A \in \mathcal{F}$ are such that $v(A_i) \geq v(B_i)$ for all $i \in N$, then $A \in \mathcal{F}_G$.*

This assumption bears an immediate interpretation as a requirement of Pareto efficiency: it asks that any provider guarantee (as captured by \mathcal{F}_G) should allow for more value to be generated whenever that is possible; if an allocation satisfies the provider guarantees, then feasible allocations where every provider can generate *more* value should also satisfy these

guarantees. Seen in this light, the assumption comes across as a natural axiomatic requirement that should be satisfied by any valid guarantees that the system might offer its providers.

In a different sense, if the set \mathcal{F}_G captures allocations that are “acceptable” to the providers (e.g., as they carry desirable guarantees), then Assumption 2 can also be interpreted as a condition on the providers’ preferences: namely, that assigning higher-value jobs to all providers should not be “less acceptable.” This suggests that the provider preferences encoded in the set of guarantees should be aligned with the value function v . The following result, which is an equivalent characterization of Assumption 2, formalizes this intuition.

Proposition 1 (\mathcal{F}_G Aligned with v). *A subset $\mathcal{F}_G \subseteq \mathcal{F}$ satisfies Assumption 2 if and only if \mathcal{F}_G satisfies*

$$\mathcal{F}_G = \arg \max_{A \in \mathcal{F}} g(A),$$

for some $g : \mathcal{F} \rightarrow \mathbb{R}$ that satisfies $g(B) \geq g(A)$ for any $A, B \in \mathcal{F}$ with $v(B_i) \geq v(A_i)$ for all $i \in N$.

A proof can be found in Appendix A. The result establishes that all valid provider guarantees correspond to value-maximizing allocations when jobs are evaluated according to some function g that preserves the same ordering as the value function v . Insofar as allocations with guarantees \mathcal{F}_G capture provider preferences, this also means that Assumption 2 exactly requires such preferences to be “aligned” with v . As we demonstrate in the next subsection, such alignments between captured value and provider guarantees (and preferences) can arise naturally in many practical settings.

2.4. Discussion and Classes of Provider Guarantees

2.4.1. Income Guarantees Under Monotonic Payment Functions.

An important class of provider guarantees satisfying Assumption 2 arises from ensuring a minimum level of total income to each provider, when every provider’s compensation is increasing in the value captured by job completions. More precisely, suppose that a provider who completes a set of jobs $S \subseteq D$ is compensated with an amount $p(S)$, where $p : \mathcal{P}(D) \rightarrow \mathbb{R}$ denotes a payment function. For a given real number τ , we define the set

$$\mathcal{F}_G(p, \tau, N) := \{A \in \mathcal{F} \mid p(A_i) \geq \tau, \forall i \in N\}, \quad (2)$$

which includes the allocations that ensure that each provider is compensated with at least an amount τ . The set $\mathcal{F}_G(p, \tau, N)$ satisfies Assumption 2 if p satisfies the property:

$$v(S) \geq v(T) \Rightarrow p(S) \geq p(T), \forall S, T \in \mathcal{P}(D). \quad (3)$$

The latter is a natural requirement for payment or compensation functions: it asks that a set of jobs carrying more intrinsic value should also command a (weakly) higher compensation when completed. An important family of payment functions that satisfy this property is proportional compensation functions of the form $p(S) = \beta v(S)$ for some $\beta \in [0, 1]$. Such proportional payments are widely used in practical revenue-sharing systems (including Upwork, Grubhub, Uber, and Lyft, among many others) where service providers retain a constant fraction of the payment made by the consumer after the job is completed. Proportional payments also include commission-based payment mechanisms that are used to compensate sales agents (see, e.g., Farley 1964, Eisenhardt 1988), as well as common bonus schemes used to incentivize employees (see, e.g., Gibbons 1998, Lazear 2000).

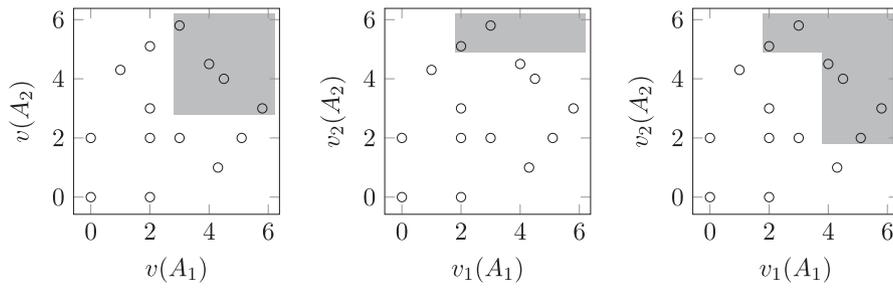
The class of minimal income guarantees could also be generalized by considering provider-specific payments and guarantees (i.e., by taking p_i or τ_i or by considering guarantees only for a subset $\hat{N} \subseteq N$ of the providers). A visual depiction of several such guarantees is shown in Figure 1, for the revenue-sharing case with $p(S) = \beta v(S)$ for $\beta = 1$. The figure shows both uniform income guarantees (with a unique τ for all providers), as well as nonuniform ones (with different τ_i for each provider). A special case of uniform income guarantees is obtained when τ is the largest value that ensures \mathcal{F}_G is nonempty; this corresponds to max-min fair allocations (see Kalai and Smorodinsky 1975 and Mas-Colell et al. 1995), which we discuss briefly in Section 2.4.3.

2.4.2. Unions and Intersections of Guarantees. Consider any collection of sets $\{\mathcal{F}_G^k\}_{k \in K}$ where each set \mathcal{F}_G^k satisfies Assumption 2. Then, $\bigcap_{k \in K} \mathcal{F}_G^k$ and $\bigcup_{k \in K} \mathcal{F}_G^k$ also satisfy Assumption 2. Considering intersections is useful for modeling restrictions that certain targeted providers find desirable or acceptable. Namely, building on the intuition developed in Proposition 1, suppose each provider $i \in N$ is endowed with a utility function u^i satisfying the comonotonicity requirements in Proposition 1.⁵ In addition, let

$$\mathcal{F}_G^i := \arg \max_{A \in \mathcal{F}} u^i(A)$$

be the allocations that maximize the provider’s utility. Then, the system could consider $\mathcal{F}_G = \bigcap_{i \in \hat{N}} \mathcal{F}_G^i$ as the allocations with guarantees. Alternatively, we could also consider a “satisficing” model (see Simon 1956) where $\mathcal{F}_G^i := \{A \in \mathcal{F} : u^i(A) \geq \tau_i\}$ are the allocations that provider i finds “acceptable” (i.e., exceeding a minimum utility threshold) and the system considers only allocations in $\bigcap_{i \in \hat{N}} \mathcal{F}_G^i$ that all providers find acceptable.

Figure 1. Examples of Provider Guarantees



Notes. Income guarantees $\mathcal{F}_G(p, \tau, N)$ for two providers ($n = 2$) compensated according to revenue-sharing agreements with $p(S) = \beta v(S)$ for $\beta = 1$. Each circle denotes a feasible allocation, with the two axes corresponding to the intrinsic values $v(A_1)$ and $v(A_2)$ for each provider. The circles in the shaded area show the revenues achievable by allocations in each \mathcal{F}_G . (Left panel) A uniform income guarantee with $\tau = 3$. (Center panel) A nonuniform income guarantee with $\tau_1 = 2, \tau_2 = 5$. (Right panel) The union of two nonuniform income guarantees with $(\tau_1, \tau_2) = (2, 5)$ and $(\tau_1, \tau_2) = (4, 2)$.

Unions of allocation sets may capture scenarios in which the system is choosing among several possible restrictions. One such example is depicted in the right panel in Figure 1. In fact, the following equivalent characterization of Assumption 2 shows that any provider guarantee satisfying it can actually be written as the union of income guarantees under monotonic payment functions.

Proposition 2. *The set of allocations with guarantees \mathcal{F}_G satisfies Assumption 2 if and only if it can be expressed as the union of income guarantees under monotonic payment functions. That is, $\mathcal{F}_G \subseteq \mathcal{F}$ satisfies Assumption 2 if and only if there exist monotonic payment functions and income guarantees $\{(p_k, \tau_k)\}_{k \in K}$ for some index set K such that $\mathcal{F}_G = \cup_{k \in K} \mathcal{F}_G(p_k, \tau_k, N)$.*

Proposition 2 implies that the class of income guarantees under monotonic payment functions is in some sense a universal generating family for all the restrictions that satisfy Assumption 2, as any such restriction can be captured by considering several alternatives from the former generating family. Although this characterization is not directly used in the proof of the following results, it adds to the understanding of the class of guarantees that satisfy Assumption 2. A proof of Proposition 2 can be found in Appendix A.

2.4.3. Fairness. Assumption 2 is also satisfied by considerations related to fairness/equity in how the jobs are allocated to providers—a common issue in settings of social justice and workforce compensation (see, e.g., Tremblay et al. 2000 and Cohen-Charash and Spector 2001). An important example arises from the broad class of α -fairness notions, first introduced by Atkinson (1970). An allocation is said to be α -fair

if it maximizes the constant elasticity social welfare function:

$$W_\alpha(\mathbf{A}) = \begin{cases} \sum_{i=1}^n \frac{v(A_i)^{1-\alpha}}{1-\alpha} & \text{for } \alpha \geq 0, \alpha \neq 1 \\ \sum_{i=1}^n \log(v(A_i)) & \text{for } \alpha = 1. \end{cases}$$

Because this welfare function is increasing in each component A_i , the associated restriction $\mathcal{F}_G = \text{argmax}_{\mathbf{A} \in \mathcal{F}} W_\alpha(\mathbf{A})$ satisfies Assumption 2 (this is an immediate corollary of Proposition 1). The constant elasticity social welfare function is concave and component-wise increasing, and thus, it exhibits diminishing marginal welfare increase as the values generated increase. In other words, if $v(A_i) < v(A_j)$, for providers i and j , then a marginal increase in $v(A_i)$ would lead to a larger welfare increase than a marginal increase in $v(A_j)$. Additionally, the rate at which marginal increases diminish is governed by the parameter α ; at higher α , there are higher incentives to increase the value allocated to providers that are worse off, which makes solutions “more fair.” This motivates the name of α as the *inequality aversion parameter*; in fact, maximizing $W_\alpha(\mathbf{A})$ for $\alpha = 0$ recovers an efficient solution, whereas $\alpha \rightarrow \infty$ recovers *max-min fair* solutions (see Kalai and Smorodinsky 1975, Mas-Colell et al. 1995).

Max-min fair allocations are inspired by the notion of Rawlsian justice (Rawls 1971), and result from uniform income guarantees under monotonic payment functions, when the guarantee τ is the largest possible value for which the restriction set \mathcal{F}_G is nonempty. More formally, we define the *max-min fair restriction under monotonic payments* as

$$\mathcal{F}_G^{\text{mm}}(p, N) := \mathcal{F}_G(p, \tau_{\text{max}}, N), \quad \text{where} \\ \tau_{\text{max}} := \max\{\tau \mid \mathcal{F}_G(p, \tau, N) \neq \emptyset\}. \quad (4)$$

3. Bounding the Value Loss Under Provider Guarantees

Our first result provides an upper bound on the relative value loss $L_\gamma(\mathcal{F}, \mathcal{F}_G)$ that holds for any feasible allocations and restrictions satisfying our assumptions.

Theorem 1 (Upper Bound on the Value Loss). *The value loss under provider guarantees is bounded above as follows:*

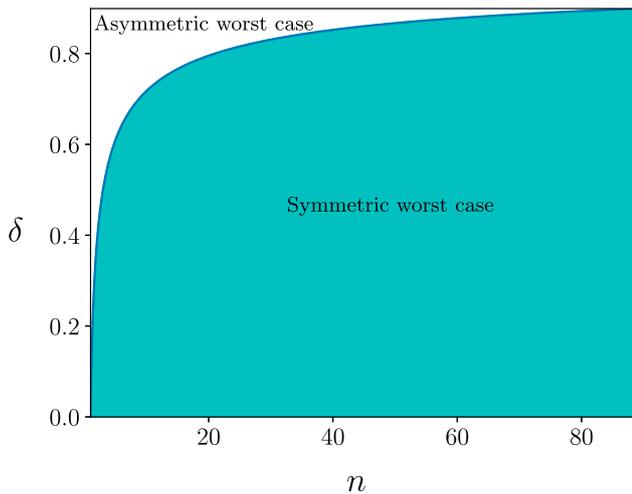
$$\sup_{\mathcal{F}, \mathcal{F}_G, \gamma} L_\gamma(\mathcal{F}, \mathcal{F}_G) \leq \max\left\{\delta, \frac{n-1}{n+(1-\delta)(n-1)}\right\}, \quad (5)$$

where the supremum is taken with respect to all $\gamma \in [\gamma_{\min}, \gamma_{\max}]^n$ and all sets $\mathcal{F}, \mathcal{F}_G$ satisfying Assumptions 1 and 2.

Theorem 1 provides a bound on the relative value loss that depends on the number of providers n and the heterogeneity level δ . Note that the bound is fully characterized by these two parameters; Figure 2 depicts the parametric regions where $\frac{n-1}{n+(1-\delta)(n-1)}$ exceeds δ (shaded area), which occurs when many providers are present (n is large) and the heterogeneity level δ is not too high. Otherwise, when the heterogeneity δ is high and/or there are only a few providers, the bound on the loss is driven solely by the heterogeneity and equals δ .

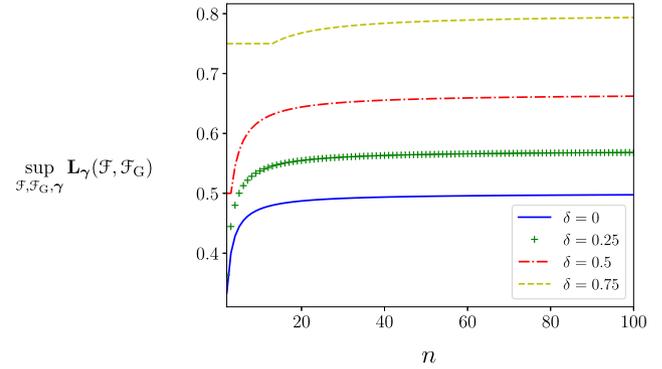
The dependency of the bound on n and δ is depicted in Figure 3. It can be seen from (5) that for any fixed heterogeneity level δ , the maximum value loss is always strictly smaller than $\frac{1}{2-\delta}$. This implies that as long as the heterogeneity is bounded, a system implementing any of the discussed provider guarantees can only incur a limited value loss, and this loss never

Figure 2. (Color online) Different Types of Worst Cases in Different Parametric Regions



Notes. Shaded area denotes values of (n, δ) such that $\delta \leq \frac{n-1}{n+(1-\delta)(n-1)}$. In the nonshaded area, the worst-case instance corresponds to asymmetric guarantees (see Instance 1), where all jobs are guaranteed to only one provider, whereas in the shaded area, the worst-case instance corresponds to a symmetric guarantee (see Instances 2 and 3), where \mathcal{F}_G corresponds to a max-min fairness guarantee.

Figure 3. (Color online) Worst-Case Loss as a Function of the Number of Providers and the Heterogeneity Level



Note. Plot of $\max\{\delta, \frac{n-1}{n+(1-\delta)(n-1)}\}$ as a function of n for different values of δ .

exceeds $\frac{1}{2}$ when providers are perfectly homogeneous ($\delta = 0$). Importantly, and as emphasized in Section 1.2, this result distinguishes our findings from many studies that document unbounded losses in contexts of price of fairness (Bertsimas et al. 2011) or price of anarchy (Koutsoupias and Papadimitriou 2009).

3.1. Main Ideas in the Proof

We defer the complete proof of Theorem 1 to Appendix A, but we briefly describe its main ideas here. First, we propose a linear programming (LP) relaxation of the problem of maximizing $L_\gamma(\mathcal{F}, \mathcal{F}_G)$ over all sets \mathcal{F} and \mathcal{F}_G , for a fixed vector γ . For this purpose, we provide a family of inequalities that connect the maximum value over the set of feasible allocations \mathcal{F} and over the set of allocations with guarantees \mathcal{F}_G . By exploiting the quasiconvexity of the optimal value of this relaxation as a function of γ , we can then maximize the loss by only considering extreme heterogeneity profiles $\gamma \in [\gamma_{\min}, \gamma_{\max}]^n$. Finally, we solve the LP relaxation for these values of γ to obtain the desired upper bound.

The next result shows that the upper bound in Theorem 1 is in fact tight, by characterizing instances and provider guarantees that achieve the worst-case loss.

Theorem 2 (Attainable Value Loss). *The bound on value loss in Theorem 1 is tight. In particular, for every δ, n , and $\varepsilon > 0$, there exist $\mathcal{F}^1, \mathcal{F}_G^1, \gamma^1$ and $\mathcal{F}^2, \mathcal{F}_G^2, \gamma^2$ satisfying Assumptions 1 and 2 so that*

$$L_{\gamma^1}(\mathcal{F}^1, \mathcal{F}_G^1) = \delta \quad (6)$$

$$L_{\gamma^2}(\mathcal{F}^2, \mathcal{F}_G^2) = \frac{n-1}{n+(1-\delta)(n-1)} - \varepsilon. \quad (7)$$

Theorem 2 states that the worst-case relative losses characterized in Theorem 1 are tight. To prove this result, we exhibit three classes of instances, where the

first one achieves the loss in (6) and the other two asymptotically achieve the loss in (7). To provide more intuition, we describe these instances for the case with $n = 2$ here and defer the general case to Appendix A.

Instance 1 (High Guarantees with High Provider Heterogeneity). Consider $n = 2$ providers with $\gamma_1 = \gamma_{\max}$ and $\gamma_2 = \gamma_{\min}$, a set of jobs D with $v(D) > 0$, a set of allocations with guarantees $\mathcal{F}_G = \{(\emptyset, D)\}$, and any feasible set \mathcal{F} with $\mathcal{F}_G \subset \mathcal{F}$. Then, $L_\gamma(\mathcal{F}, \mathcal{F}_G) = \frac{\gamma_{\max} - \gamma_{\min}}{\gamma_{\max}} = \delta$.

In Instance 1, the system has a set of jobs D to assign to two providers, and it is possible to assign all the jobs to a single provider.⁶ The first provider generates more value for the system ($\gamma_1 \geq \gamma_2$), so the value-maximizing allocation would be (D, \emptyset) , assigning all jobs to that provider. However, the only allocation with guarantees is (\emptyset, D) , requiring all jobs to be assigned to the second provider, and thus generating a loss because of heterogeneity. This results in worst-case losses when heterogeneity is high and the number of providers is not too large, as depicted in Figure 2.

The high losses in Instance 1 are enabled by the joint presence of two key features: a high degree of heterogeneity in the providers' effectiveness to generate value (with one provider maximally effective, and another minimally effective) and asymmetric guarantees that allocate all the jobs to the least effective provider. This asymmetry is critical: if the guarantees had been "symmetrized" (e.g., by considering the union of all permutations of which provider receives the full set of jobs), then the loss would vanish. So for heterogeneity to be a critical driving force, it must be accompanied by asymmetric guarantees that are misaligned in order to force jobs being assigned to less effective providers. The inefficiency in Instance 1 can thus play a particularly prominent role in service platforms that may favor new providers by guaranteeing them a higher workload, when such providers also have less experience and thus generate less value, for example, in terms of productivity as well as goodwill gain and customer satisfaction.⁷

The following instance achieves the bound on the relative loss that is given in (7).

Instance 2 (Max-Min Fairness with Monotonic Payments). Consider $n = 2$ providers with $\gamma_1 = \gamma_{\max} \geq \gamma_2 = \gamma_{\min}$, and a set of jobs $D = \{d_1, d_2, d_3\}$ with $v(d_1) = v(d_2) = 1$ and $v(d_3) = 1 - \kappa$ for some $\kappa > 0$. The feasible allocations \mathcal{F} are depicted in Figure 4: jobs can be assigned together only if the corresponding segments are nonoverlapping. The allocations with guarantees are given by all max-min fair allocations $\mathcal{F}_G^{\text{MM}}(p, N)$ under any strict monotonic payment function p , as described in (4). Then, the value-maximizing allocation

is $(\{d_1, d_2\}, \{d_3\})$, the value-maximizing allocation with guarantees is $(\{d_1\}, \{d_2\})$,⁸ and the relative value loss is

$$L_\gamma(\mathcal{F}, \mathcal{F}_G) = \frac{(2\gamma_{\max} + (1 - \kappa)\gamma_{\min}) - (\gamma_{\max} + \gamma_{\min})}{2\gamma_{\max} + (1 - \kappa)\gamma_{\min}}.$$

In particular, for any $\epsilon > 0$ there exists a $\kappa > 0$ such that $L_\gamma(\mathcal{F}, \mathcal{F}_G) = \frac{1}{3 - \delta} - \epsilon$, and thus,

$$L_\gamma(\mathcal{F}, \mathcal{F}_G) \xrightarrow{\kappa \rightarrow 0} \frac{1}{3 - \delta} = \frac{n - 1}{n + (1 - \delta)(n - 1)}.$$

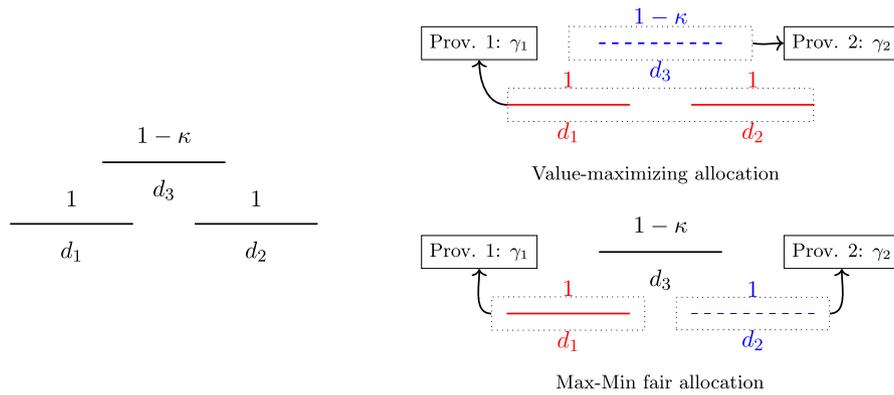
Instance 2 describes a situation where the system allocates jobs that have a certain time duration and where two jobs that overlap in time cannot be assigned to the same provider—a situation that occurs routinely in ride-sharing platforms such as Uber or Lyft. A distinctive property of Instance 2 is that the system can either assign all the high-value jobs d_1, d_2 to one provider while assigning the low-value job d_3 to the other provider, or it can distribute the two high-value jobs among the providers and *not* allocate the low-value job. The value loss is thus created because the guarantee imposes the latter allocation, which results in unassigned jobs, and this loss increases as the unassigned job d_3 is very close in value to each of the allocated jobs d_1, d_2 .

Instance 2 thus showcases two new drivers for the value loss. First is the structure of the set of feasible allocations \mathcal{F} , which contains certain exclusion constraints that force job d_3 to be incompatible with both jobs d_1 and d_2 . Note that if job d_3 could be allocated together with either of the two other jobs, then the loss $L_\gamma(\mathcal{F}, \mathcal{F}_G)$ would vanish. We revisit such exclusion constraints in Section 4, where we analyze their impact on the value loss in more detail. Second is the variation in the intrinsic value of jobs; although the loss grows as jobs become more similar in value (i.e., as $\kappa \rightarrow 0$), some variation is in fact critical: if $\kappa = 0$, the loss $L_\gamma(\mathcal{F}, \mathcal{F}_G)$ would again vanish. Our next instance shows that this variation is actually *not* critical to achieving a worst-case loss.

Instance 3 (Max-Min Fairness with Monotonic Payments and Equal-Valued Jobs). Consider $n = 2$ providers with $\gamma_1 = \gamma_{\max}$ and $\gamma_2 = \gamma_{\min}$, and a set of jobs $D = \{d_1, d_2, d_3, d_4, d_5\}$ with $v(d_i) = 1$ for each $i \in N$. The feasible allocations are depicted in Figure 5: jobs can be assigned together only if the corresponding segments are nonoverlapping. The allocations with guarantees correspond to the max-min fair allocations under any strict monotonic payment function, as described in (4). Then, the value-maximizing allocation is $(\{d_1, d_2, d_3, d_4\}, \{d_5\})$, and a value-maximizing allocation with guarantees is $(\{d_1, d_2\}, \{d_3, d_4\})$.⁹ Therefore,

$$L_\gamma(\mathcal{F}, \mathcal{F}_G) = \frac{4\gamma_{\max} + \gamma_{\min} - 2(\gamma_{\max} + \gamma_{\min})}{(4\gamma_{\max} + \gamma_{\min})}.$$

Figure 4. (Color online) Symmetric Worst-Case Instance



Notes. The left panel depicts Instance 2 that consists of $n = 2$ providers and three jobs $D = \{d_1, d_2, d_3\}$ represented by segments. Jobs can be assigned together if the corresponding segments are nonoverlapping. The right panel depicts the value-maximizing allocation at the top, and the Max-Min fair allocation at the bottom.

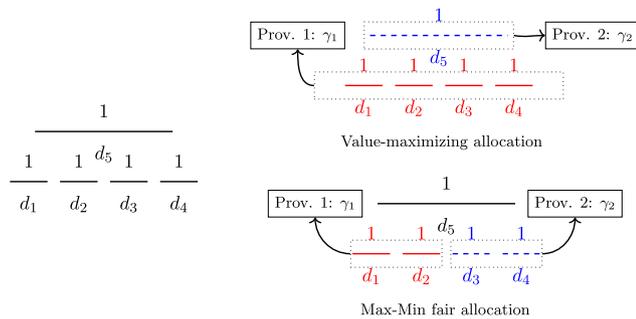
Instance 3 can be generalized to any given number of providers n , by taking a set of $(t + 1)n + t$ jobs for any integer $t > 0$ and keeping the same structure of \mathcal{F} (details on the construction are provided in Instance A.3 of Appendix A). This generalization yields a worst-case loss of

$$\begin{aligned} L_{\gamma}(\mathcal{F}, \mathcal{F}_G) &= \frac{(t + 1)n\gamma_{\max} + t(n - 1)\gamma_{\min} - (t + 1)(\gamma_{\max} + (n - 1)\gamma_{\min})}{(t + 1)n\gamma_{\max} + t(n - 1)\gamma_{\min}} \\ &= \frac{t(n - 1) + \delta(n - 1)}{t(n + (n - 1)(1 - \delta)) + n}. \end{aligned}$$

Therefore, for any $\epsilon > 0$, there exists a t large enough such that $L_{\gamma}(\mathcal{F}, \mathcal{F}_G) = \frac{n-1}{n+(1-\delta)(n-1)} - \epsilon$.

Instance 3 shares certain similarities with Instance 2. Both rely on having a large number of providers and jobs, and on the presence of exclusion constraints in the set of feasible allocations \mathcal{F} that prevent certain jobs from being allocated together. Also, the worst-case

Figure 5. (Color online) Symmetric Worst-Case Instance with Equal-Valued Jobs



Notes. The left panel depicts Instance 3: there are $n = 2$ providers and five jobs $D = \{d_1, d_2, d_3, d_4, d_5\}$ represented by segments. Jobs can be assigned together if the corresponding segments are nonoverlapping. The right panel depicts the value-maximizing allocation at the top and the value-maximizing max-min fair allocation at the bottom.

guarantees in both instances correspond to symmetric max-min fairness guarantees obtained under any strictly monotonic payment function. This suggests that in allocation systems with many providers and sufficiently “standardized” jobs (so that provider heterogeneity is not too large), symmetric fairness guarantees can be critical drivers of loss, as depicted in Figure 2. Lastly, it is worth emphasizing that although both Instances 2 and 3 involve some heterogeneity in the providers’ effectiveness, this is not critical to identifying these as worst-case instances; in fact, both instances yield worst-case losses when providers are homogeneous ($\delta = 0$).

Instance 3 also exhibits some notable differences from Instance 2: it relies on jobs with identical value, and it requires an arbitrarily large number of jobs to achieve the worst-case loss. We return to discuss each of these key drivers—the variation in job values and the supply/demand imbalance—and their relationship in more detail in Section 4. Additionally, note that the value of jobs in Instance 3 cannot be proportional to the size of the associated intervals; in contrast, Instance 2 can be constructed to satisfy this, which brings it closer to several practical applications, such as when intervals represent the time to complete a job and the value generated is proportional to this time.

We conclude our discussion by noting that all the loss drivers that are critical in Instances 2 and 3 are irrelevant in Instance 1, where the loss is entirely driven by the heterogeneity in provider effectiveness. To see this, note that all jobs in Instance 1 are assigned under all the allocations with guarantees \mathcal{F}_G ; thus, no value is lost because of jobs that remain unallocated as a result of physical constraints or restrictions imposed by the provider guarantees. Additionally, note that as long as jobs have some positive intrinsic value in Instance 1, any distribution of values could lead to

worst-case losses; thus, the precise job values are completely irrelevant. Lastly, note that Instance 1 can be readily generalized to obtain the same loss in a setting with an arbitrary number of providers, as long as at least one maximally effective and one minimally effective provider exist and the guarantees require allocating jobs only to the latter (see Appendix A); thus, having a large number of providers or imbalanced supply/demand does not impact the loss in Instance 1.

4. Analysis of Key Loss Drivers

The previous section highlighted several potential drivers for value loss. Perhaps the most prominent of these is provider heterogeneity: when providers differ in their effectiveness to generate value (i.e., δ is large), this heterogeneity becomes the dominant loss driver, as evidenced in Instance 1. In this section, we focus our discussion on several additional loss drivers: (i) the structure of the set of feasible allocations \mathcal{F} ; (ii) the variation in the intrinsic values of jobs; (iii) the balance between supply (number of providers) and demand (number of jobs); and (iv) the integrality of allocations.

4.1. Structural Properties of the Set of Feasible Allocations

When heterogeneity is not the main loss driver, the structure of the set of feasible allocations \mathcal{F} can be crucial, as Instances 2 and 3 demonstrated; recall that each of those instances involved certain exclusion constraints, whereby some jobs could not be allocated together to the same provider. Our next example shows that such constraints are critical: when the set \mathcal{F} only includes constraints on how many jobs can be allocated together but without other explicit exclusion constraints, the value loss vanishes. Notice how we assume that all providers are homogeneous (all γ_i are equal) in order to eliminate the effect of heterogeneity on the value loss.

Instance 4. Consider a case with homogeneous providers, $\gamma_i = \gamma$, $\forall i \in N$. Let $\{k_i\}_{i=1}^n$ be n positive integers, and suppose that $\mathcal{F} = \{(A_1, \dots, A_n) \mid A_i \subseteq D, |A_i| \leq k_i, \forall i, \text{ and } A_i \cap A_j = \emptyset, \forall i \neq j\}$.

Proposition 3. Instance 4 satisfies $L_\gamma(\mathcal{F}, \mathcal{F}_G) = 0$ for any set of allocations with guarantees \mathcal{F}_G .

This result becomes even more striking when the structure of the feasible sets in Instances 2 and 3 is further broken down. In particular, note that the set of jobs in each of those instances is composed from some jobs that can be allocated together (with no further constraints) and a single job whose allocation precludes a provider from executing any other job. Our next instance generalizes these structures.

Instance 5. Consider any γ , a set of jobs $D = S \cup C$, and a set of feasible allocations of the form:

$$\mathcal{F} = \{(A_1, \dots, A_n) \mid A_i \cap A_j = \emptyset, \forall i \neq j \text{ and for every } i \in N, \text{ either } A_i \cap C = \emptyset \text{ or } |(A_i \cap C)| = 1 \text{ and } A_i \cap S = \emptyset\}.$$

The jobs described in Instance 5 can be divided into a set of unconstrained jobs S and a set of jobs C , each of which cannot be allocated together with any other job in D . Each of these sets of jobs considered *in isolation* would give rise to a set of feasible allocations that would conform to the premises of Instance 4 (with capacity $k_i = \infty$ or $k_i = 1$, respectively), and thus result in zero loss. It is thus striking that by simply combining these two sets as done in Instance 5, one in fact obtains an instance that could either have zero loss or a worst-case loss, as formalized in the next result.

Proposition 4. Consider Instance 5, then

- i. If $v(d) = v$, for all $d \in D$, and $|S| < 2n$, $L_\gamma(\mathcal{F}, \mathcal{F}_G) = 0$ for any uniform income guarantee \mathcal{F}_G .
- ii. If $v(d) = 1 - \kappa$ for all $d \in C$, $v(d) = 1$ for all $d \in S$, $|C| = n - 1$, and $|S| = n$, we recover Instance 2 by taking \mathcal{F}_G as the max-min fair allocations under any strictly monotonic payment function.

Proposition 4 implies that the feasible set structure can carry significant impact, but also that this structure in isolation is not a good predictor of value loss: the same structure could generate very large or very small losses, depending on other problem features such as the value of jobs or the imbalance between supply (number of providers) and demand (number of jobs). Notice how all conditions in (ii) satisfy the conditions in (i), except for the deviation in the intrinsic values of jobs in C . This also shows how small changes in an instance can lead to large changes in the value loss.

We conclude by noting that the symmetry of the set of feasible allocations \mathcal{F} is also very important for our results. If this assumption was relaxed (e.g., if providers had different sets of jobs they could complete), then we can actually achieve a worst-case loss that asymptotically approaches 100% as the number of agents grows large (see Instance A.4 of Appendix A). This instance is actually inspired by the bandwidth allocation problem considered in Bertsimas et al. (2011) and matches the upper bound on the price of fairness proved therein. This finding is consistent with our earlier result concerning the impact of heterogeneity in the providers' effectiveness to generate value from jobs, and it reinforces the idea that when providers are sufficiently "different," this can critically drive the losses, making them unbounded in extreme cases. To further explore this, in Appendix B we analyze numerically the average relative loss for a

family of instances with a varying degree of *symmetry*, measured by the proportion of all the jobs that each provider can perform. Interestingly, we find that as the instances become more asymmetric, the average losses decrease.

Intuitively, when each provider can perform a smaller fraction of all jobs, the probability of two providers being able to perform the same job decreases. In the extreme case this leads to zero loss: if no job can be performed by two providers, then the allocation problems can be separated into disjoint problems for each provider, which by Proposition 5, implies that the value loss will be zero. Nevertheless, this decrease in the average loss is tied to the assumption that the subset of jobs a provider can perform is independent across providers. In fact, in Instance A.4 of Appendix A, the high loss is driven by a specific segmentation of the providers into two groups, each with a particular set of jobs they can perform. This suggests that the relative value loss is significantly influenced by the specific notion of symmetry considered and that the impact of symmetry can be nontrivial: allowing for very high asymmetry might lead to very high worst-case losses, but asymmetry might also decrease the average losses in some cases.

4.2. Variation in Intrinsic Job Values

Instance 2 showed that the difference in the intrinsic value of jobs can be a key driver of the value loss and that this loss can decrease as jobs have increasingly different intrinsic values. To further explore the impact of this feature, we now consider a slight modification of Instance 2 where we introduce a random perturbation in the value of one of the jobs, and we consider the expected loss as a function of the variance of this perturbation. Because each realization of the random perturbation corresponds to a particular Instance 2, when the variance of the perturbation is high we will obtain higher differences in the intrinsic values of jobs.

Example 1. Consider Instance 2 with $\gamma_{\min} = \gamma_{\max}$, and assume that κ is uniformly distributed $\kappa \sim U[-\frac{\Delta}{2}, \frac{\Delta}{2}]$, so that $v(d_3) \sim U[1 - \frac{\Delta}{2}, 1 + \frac{\Delta}{2}]$. By taking the expectation of the loss with respect to this random variable, we get

$$\begin{aligned} \mathbb{E}[\mathbf{L}_\gamma(\mathcal{F}, \mathcal{F}_G)] &= \mathbb{E}\left[\frac{v(d_3)}{2 + v(d_3)} \mathbb{1}_{\{v(d_3) \in [0, 1]\}}\right] \\ &= \int_{\max\{0, 1 - \frac{\Delta}{2}\}}^1 \frac{1}{\Delta} \left(\frac{s}{2 + s}\right) ds \\ &= \frac{1}{\Delta} \left(1 - 2 \log(3) - \max\left\{0, 1 - \frac{\Delta}{2}\right\}\right) \\ &\quad + 2 \log\left(2 + \max\left\{0, \frac{\Delta}{2}\right\}\right). \end{aligned}$$

Note that $\mathbb{E}[\mathbf{L}_\gamma(\mathcal{F}, \mathcal{F}_G)]$ is decreasing in Δ , and because the variance of $v(d_3)$ equals the variance of κ , which equals $\frac{\Delta^2}{12}$, this implies that $\mathbb{E}[\mathbf{L}_\gamma(\mathcal{F}, \mathcal{F}_G)]$ is decreasing in the variance of $v(d_3)$: higher variance implies lower expected loss.

Interestingly, this example suggests that when intrinsic values are randomly generated, higher variance in values could actually *reduce* the impact of implementing provider guarantees, and result in a *lower* expected loss. The example can be generalized to a case with n providers (see Example A.1 in Appendix A), and we also confirm its robustness in a more realistic setting as part of our numerical exercise in Section 5. However, it is also worth noting that although these examples suggest that lower variation tends to lead to higher value losses, the relationship may exhibit a sharp discontinuity when there is no variation at all. This is already evident in Instance 2, where requiring all jobs to have the same identical value reduces the relative loss under any uniform income guarantees to zero; this pattern continues to occur in many of the data-driven instances we analyze in Section 5.

4.3. Supply-Demand Balance

Instance 3 suggested that the imbalance between supply (number of providers) and demand (number of jobs) can also critically drive the value loss. To build further intuition for this, we examine certain extreme cases that allow an analytical characterization. The next result shows that when a single provider or a sufficiently large number of providers is present, the value loss vanishes. We again limit the effect of heterogeneity on the value loss by assuming that all providers are homogeneous.

Proposition 5. Assume that $\gamma_{\min} = \gamma_{\max}$, and consider any set of allocations with guarantees \mathcal{F}_G . Then, $\mathbf{L}_\gamma(\mathcal{F}, \mathcal{F}_G) = 0$ if either $n = 1$ or $n \geq |D|$.

Additionally, if we restrict attention to uniform income guarantees and jobs with identical intrinsic values, then the value loss would be zero for an even larger number of providers, as formalized by the following result.

Proposition 6. Assume that $\gamma_{\min} = \gamma_{\max}$, and let $n > |D|/2$ and $v(d) = v$ for all $d \in D$. Then, $\mathbf{L}_\gamma(\mathcal{F}, \mathcal{F}_G) = 0$, for any uniform income guarantee \mathcal{F}_G .

The intuition behind this result is that when there are insufficient jobs to guarantee each provider at least two jobs, then there will be no value loss from guaranteeing the same level of income to all providers. This suggests that the value loss is likely small when the supply-demand imbalance is high (i.e., the ratio of supply to demand is either very low or very high). Although it is hard to analytically prove this

more generally, we confirm it in our numerical tests in Section 5, where we find that increasing the number of providers for a fixed number of jobs initially increases and eventually decreases the value loss, on average.

4.4. Integrality of Allocations

The integrality of allocations is critical for the appearance of loss in our framework: when partial allocations of jobs are possible, the loss vanishes for any uniform income guarantee \mathcal{F}_G . To formalize this, we first define the set of fractional allocations \mathcal{F}^c obtained by allowing an arbitrary mixing of allocations from \mathcal{F} :

$$\mathcal{F}^c = \left\{ \left(\{\theta_j\}_{j=1}^k, \{A^j\}_{j=1}^k \right) \mid 0 \leq \theta_j \leq 1, \sum_{j=1}^k \theta_j = 1, A^j \in \mathcal{F}, \forall j \in \{1, \dots, k\}, k \geq 0 \right\}. \quad (8)$$

Each tuple of \mathcal{F}^c represents a specific mixing of allocations from \mathcal{F} and can be interpreted as allocating a fraction θ_j of the jobs from each allocation A^j . Hence, for $C = (\{\theta_j\}_{j=1}^k, \{A^j\}_{j=1}^k) \in \mathcal{F}^c$, let us denote by $C_i = (\{\theta_j\}_{j=1}^k, \{A_i^j\}_{j=1}^k)$ the specific mixing allocated to provider i , and let us extend our value functions such that

$$v(C_i) = \sum_{j=1}^k \theta_j v(A_i^j). \quad (9)$$

Additionally, given any payment function $p(\cdot)$ that is monotonic with respect to $v(C_i)$, we define the set of uniform income guarantees $\mathcal{F}_G^c = \{C \in \mathcal{F}^c \mid p(C_i) \geq \tau, \text{ for } i \in N\}$.

The following result shows that the relative loss vanishes for any uniform income guarantees.

Proposition 7. *Given any γ and any set of feasible allocations \mathcal{F} , consider the set \mathcal{F}^c defined in (8) and the extension of $v(C)$ defined in (9). Then, $\mathbf{L}_\gamma(\mathcal{F}^c, \mathcal{F}_G^c) = 0$ for any set of uniform income guarantees under monotonic payments \mathcal{F}_G^c .*

The intuition behind Proposition 7 is that when fractional allocations are possible, the system can simply consider an allocation obtained by mixing with an equal proportion $\frac{1}{n!}$ all the permutations of a particular value-maximizing allocation. This new allocation would still achieve the maximum value while also allowing each agent to generate exactly the same value, and thus remaining feasible under any uniform income guarantee.

Our definition of fractional allocations is motivated by settings with finite periods and similar jobs being allocated in each period. Thus, although the jobs being allocated cannot be partitioned in each period, overall the allocations may be randomized across periods. Therefore, our result in Proposition 7 suggests that for certain provider guarantees defined with respect to average monotonic payment functions,

the loss can be driven close to zero by correctly randomizing the allocations.

5. Numerical Analysis of Real-World and Synthetic Data

To demonstrate the impact of our findings in a practical context, we next provide a numerical study that is based on real and synthetic data.

5.1. Basic Setup

We design our study around the problem of allocating requests for transportation (taxi rides) to drivers. The set of service providers consists of n drivers, where we consider values of $n \in \{2, 3, \dots, 30\}$. We assume that drivers are homogeneous, and thus, $\gamma_{\min} = \gamma_{\max} = \gamma$, which we normalize without loss of generality to one. The set of jobs D corresponds to trip requests that arrived in a particular time window. Each job/trip is specified as a continuous time interval given by a start time and a trip duration. Two trips can be allocated together to the same driver only if the corresponding time intervals do not overlap. We therefore construct the set of feasible allocations \mathcal{F} by putting together all the allocations (A_1, \dots, A_n) consisting of n mutually exclusive subsets of trips ($A_i \subseteq D, A_i \cap A_j = \emptyset, \forall i \neq j$) where each subset A_i contains only nonoverlapping trips. The set of allocations with guarantees \mathcal{F}_G is obtained by considering max-min fairness considerations under revenue sharing, as discussed in Section 2.4.1. To obtain the relative value loss $\mathbf{L}_\gamma(\mathcal{F}, \mathcal{F}_G)$, we compute the value-maximizing allocation and the best restricted allocation by solving integer programming problems when allocations are picked from \mathcal{F} and \mathcal{F}_G , respectively. Further details on our setup are provided in Appendix B.

5.2. Real Data

We generate problem instances using a publicly available data set containing all the completed taxi trips in New York City (NYC Taxi and Limousine Commission 2016) for January 2016. The record for every completed trip includes the total fare paid, the starting and ending locations, the starting time, and the trip duration. Each set of jobs we considered consists of trip requests originating and ending in a specific neighborhood of the city; we considered in separation Midtown Manhattan, Upper West Side, or Upper East Side. Limiting the geographical area ensures that part (ii) of Assumption 1 concerning the standardized nature of jobs is satisfied in that any provider in the area can perform any subset of (non-overlapping) trips; in Appendix B, we present more complex feasibility restrictions that allow us to consider all trips in the city while still satisfying this assumption. For each neighborhood, we focus on

the first week of January 2016; we consider, for each particular day in that week, all the trips completed between 9 a.m. and 5 p.m., a time segment during which the number of trips per minute was approximately constant. We partition the time horizon into intervals of w minutes each, where we considered $w \in \{10, \dots, 20\}$. For each of these intervals, we sampled uniformly 30 trips to generate a problem instance. We consider the value generated by each trip as the total trip fare that was paid to the driver.

5.3. Synthetic Data

To better control the impact of different parameters on the value loss, we also construct synthetic instances. We obtain these by first considering a particular time window $(0, x] \subseteq \mathbb{R}$ for different values of $x \in (1, 3)$. We then generate 30 trips as subintervals of $(0, x]$, with the starting point of the trip sampled uniformly, and the trip length drawn from a truncated normal distribution; we use a mean of one and several coefficients of variation $cv \in (0, 0.6]$ for the duration. We fix the value produced by each job as the length of the associated interval.

5.4. Results and Discussion

We next present a brief summary of the numerical findings, and we direct the reader to Appendix B for a more complete analysis. Figures 6 and 7 depict specific examples of feasible sets \mathcal{F} generated by the data-driven instances, together with the optimal allocations with and without the max-min fair guarantees. In these graphs, each vertex represents a trip and has a label that corresponds to the trip value; two trips are joined by an edge if they overlap in time, and an allocation is a (possibly partial) coloring of the graph

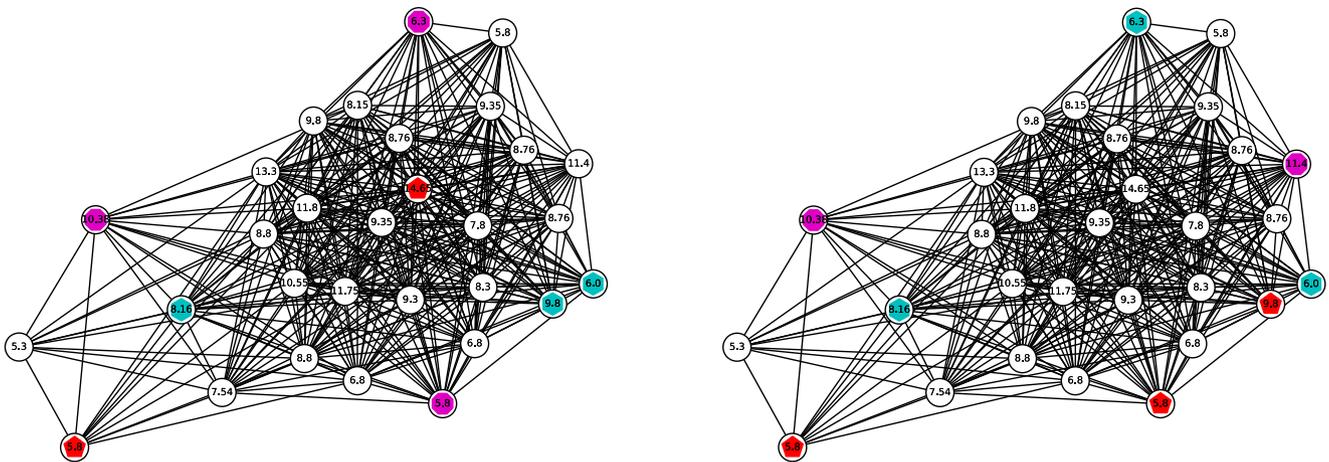
with n colors. In the instance depicted in Figure 7, all values were taken to be equal. This structure resembles Instance 5, in that there is a relatively large set of jobs that are mutually exclusive combined with smaller sets of jobs that can be allocated together.

In all our instances, when all the job values are set equal—so that the variation in values disappears—we obtain a loss of zero. This is consistent with Proposition 4, and is reasonable to expect precisely because many of our data-driven instances match Instance 5, which is the premise for Proposition 4 (refer again to Figures 6 and 7 and our earlier discussion). This implies that in the instances that we studied numerically, the variation in job values is a prominent driver of loss.

In Figure 8, we provide a representative example of the results we obtain. Both panels depict the relative value loss as a function of the number of providers. The left panel shows the average and maximum loss in the instances generated from real-world data, and the right panel corresponds to the average loss in the synthetic instances, for different coefficients of variation. In the taxi data that we considered, the coefficient of variation in job values was 0.48, so that the magnitude of the losses is consistent in the two examples. Moreover, that the maximum and the average loss are relatively close in the left panel suggests that losses come from structural properties of the instances rather than a low-frequency occurrence of instances with high relative value loss.

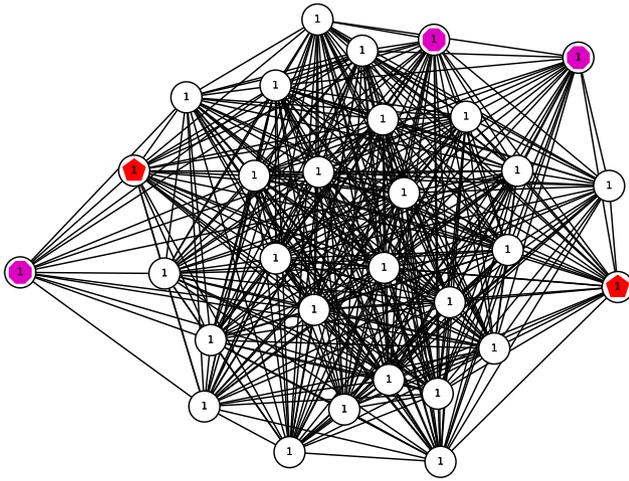
Both charts confirm several of our earlier observations. The right panel in Figure 8 shows that the maximum loss decreases with the variability in job values, which is consistent with our discussion in Example 1. Additionally, the charts are consistent

Figure 6. (Color online) Graph Representation of the Set of Feasible Allocations \mathcal{F} of an Instance



Notes. Each node represents a trip, labels represent the fare, and two nodes are connected by an edge when they cannot be allocated together. (Left panel) A value-maximizing allocation of the jobs to $n = 3$ providers, with the allocation to each provider given by a different color and shape. The total value allocated is \$66.89, and the provider with the smallest allocation receives \$20.45. (Right panel) A (value-maximizing) max-min fair allocation. The total value allocated is \$63.64, and the provider with the smallest allocation receives \$20.46.

Figure 7. (Color online) Graph Representation of \mathcal{F} for an Instance Where Trips Have Identical Intrinsic Values



Notes. Each node represents a trip, and two nodes are connected by an edge when they cannot be allocated together. A value-maximizing and max-min fair allocation with two providers is represented by the coloring and shapes in the nodes. The maximum number of trips that can be allocated to two providers is five, by allocating the only three trips that can be completed together to one provider, and two other trips to the other provider.

with the results in Proposition 5 and highlight the same qualitative features. For example, the value loss has a unique peak that appears for a ratio of supply to demand between $\frac{2}{3}$ and $\frac{1}{2}$, and decreases to 0 when the number of providers becomes sufficiently high or sufficiently low.

In addition to confirming several of our analytical findings, these numerical results also imply that the value loss associated with implementing provider guarantees in particular settings may be significantly smaller

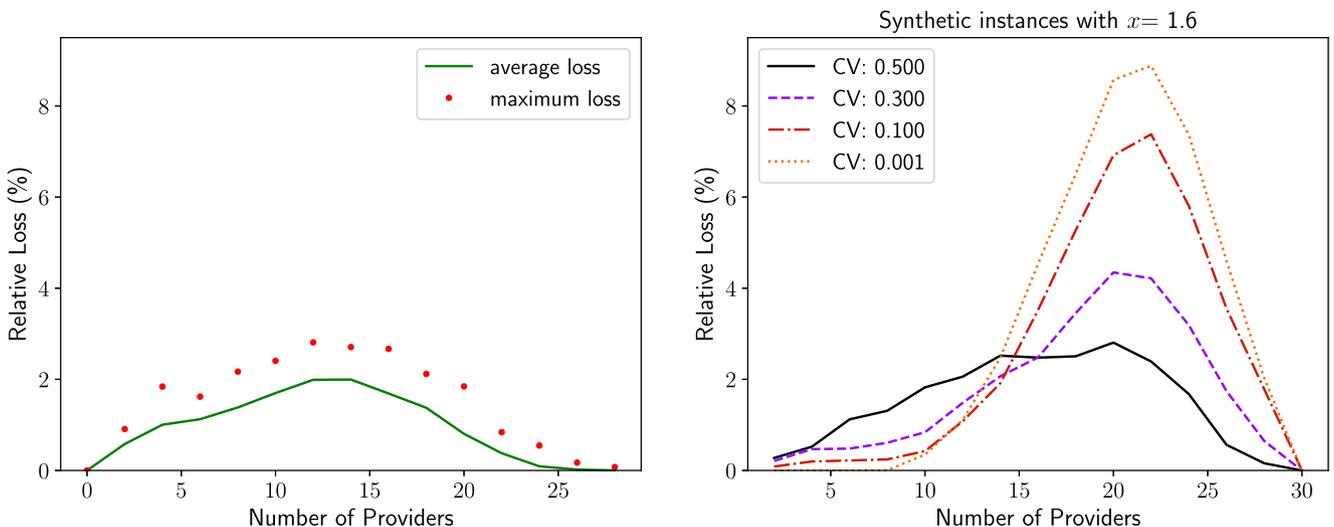
than the worst-case value loss that was characterized in Theorem 1: the relative loss did not exceed 10% in the instances generated synthetically and did not exceed 4% in the instances generated from real-world data. Together with the rest of our numerical findings, this suggests that the exact (relative) losses may be significantly smaller in particular settings of practical interest.

6. Conclusions, Limitations, and Future Research

In this paper, we established that the relative value loss because of a broad class of provider guarantees is bounded. We further showed that the worst-case losses are primarily driven by fairness considerations when many providers are present, and by the heterogeneity in the providers’ effectiveness to generate value when fewer providers are present. We analyzed several additional loss drivers, finding that both a high variation in the intrinsic values of the jobs as well as a very imbalanced (i.e., either very high or very low) ratio of supply to demand would lead to *smaller* losses. Finally, we confirmed several of the findings numerically using both real and synthetic data and documented that value losses in specific practical settings may be significantly lower than the (theoretical) worst-case values.

These findings and certain limitations in our modeling framework motivate future work of both theoretical and practical nature. Having established that losses are bounded in a general setting and under a broad class of provider guarantees, one could also seek to quantify losses in more specific settings or for subclasses of guarantees obtained as special cases of our framework. For instance, one could seek a parametrized upper bound on the relative loss when guaranteeing a specific

Figure 8. (Color online) Relative Value Loss in Particular Instances



Notes. (Left panel) Average and maximum values of the relative value loss $L_r(\mathcal{F}, \mathcal{F}_G)$ as a function of the number of providers, for instances with 30 jobs constructed from the data, using a time window of $w = 16$ minutes. (Right panel) Average value of the relative value loss $L_r(\mathcal{F}, \mathcal{F}_G)$ as a function of the number of providers for synthetic instances with different coefficients of variation (CVs), and using the parametric value $x = 1.6$.

income level to providers, or to quantify losses in specific operational settings such as ride-sharing platforms for instance, which would be closer in spirit to our numerical exercise in Section 5. These experiments also suggest that it may be worth studying the average-case loss as opposed to the worst-case loss, both because of the substantial gap between the two and also because certain qualitative effects may be provable under the former setting but not the latter; for instance, the average relative loss seems to be a unimodal function of the number of providers, whereas the worst-case loss does not.

Several of our results also suggested that the degree of symmetry in a set of feasible allocations can play a critical and nontrivial role: although asymmetry could in principle lead to unbounded worst-case losses, allowing some asymmetry could actually help reduce empirical losses in some cases of practical interest. To that end, it would be very interesting to understand more deeply the role of symmetry (e.g., by relaxing our standing Assumption 1(ii) and introducing a tractable notion that allows a bounded degree of asymmetry in the providers' capability to execute jobs).

Lastly, and from a more prescriptive viewpoint, losses remaining bounded for many types of guarantees also open the path to exploring policies that could achieve these guarantees dynamically and in an online fashion under partial information, when the streams of future jobs are unknown.

Acknowledgments

The authors thank the associate editor and three anonymous reviewers for their helpful comments that helped improve the quality of the paper.

Appendix A. Proofs and Examples

For ease of notation, let $\mathcal{F}^* \stackrel{\text{def}}{=} \arg \max_{A \in \mathcal{F}} \sum_{i=1}^n \gamma_i v(A_i)$, and $\mathcal{F}_G^* \stackrel{\text{def}}{=} \arg \max_{B \in \mathcal{F}_G} \sum_{i=1}^n \gamma_i v(B_i)$.

We now prove Propositions 1 and 2.

Proof of Proposition 1. If we have that $\mathcal{F}_G = \arg \max_{A \in \mathcal{F}} g(A)$, then if $A \in \mathcal{F}_G$, and $B \in \mathcal{F}$ satisfy that $v(B_i) \geq v(A_i)$, for each i , then we know by the statement of the proposition that $g(B) \geq g(A)$, which implies that $B \in \arg \max_{A \in \mathcal{F}} g(A) = \mathcal{F}_G$, proving that Assumption 2 holds. On the other hand, if we assume that Assumption 2 holds for a certain \mathcal{F}_G , then let us consider the function $g : \mathcal{F} \rightarrow \mathbb{R}$, defined by $g(A) = \mathbb{1}_{\mathcal{F}_G}(A)$, that takes the value 1 when $A \in \mathcal{F}_G$, and 0 otherwise. Thus, by definition $\mathcal{F}_G = \arg \max_{A \in \mathcal{F}} g(A)$. Moreover, if $g(A) = 1$, and $B \in \mathcal{F}$ is such that $v(B_i) \geq v(A_i)$, for all $1 \leq i \leq n$, then, because \mathcal{F}_G satisfies Assumption 2, $B \in \mathcal{F}_G$, which implies that $g(B) = 1 \geq g(A)$. Therefore, $g(\cdot)$ satisfies the condition of Proposition 1, which concludes the proof. \square

Proof of Proposition 2. It is clear that if we have a union of income guarantees, then Assumption 2 is satisfied; thus, we only need to prove that any \mathcal{F}_G that satisfies this assumption can be expressed as such union. For this, take any \mathcal{F}_G that satisfies Assumption 2, and consider for each $A \in \mathcal{F}_G$, the guarantee $\mathcal{F}_G^A = \{B \in \mathcal{F} \mid v(B_i) \geq v(A_i), \forall i \in N\}$. We claim then $\cup_{A \in \mathcal{F}_G} \mathcal{F}_G^A = \mathcal{F}_G$. Clearly, $\mathcal{F}_G \subseteq \cup_{A \in \mathcal{F}_G} \mathcal{F}_G^A$ because each $A \in \mathcal{F}_G^A$. Moreover, if $C \in \cup_{A \in \mathcal{F}_G} \mathcal{F}_G^A$, then there exists $A \in \mathcal{F}_G$, such that $v(C_i) \geq v(A_i), \forall i \in N$, but then by Assumption 2, $C \in \mathcal{F}_G$, which concludes the proof. \square

Now we will prove Theorems 1 and 2. We begin by proving an optimality condition that any allocations $A \in \mathcal{F}^*$ and $B \in \mathcal{F}_G^*$ must satisfy.

Lemma A.1. For any fixed \mathcal{F} and \mathcal{F}_G , let $A \in \mathcal{F}^*$, and $B \in \mathcal{F}_G^*$, then

$$v(A_i \setminus (B_1 \cup \dots \cup B_n)) \leq v(B_j \setminus A_i), \quad \forall i, j \in N. \quad (\text{A.1})$$

Proof. Assume by contradiction that $v(A_i \setminus (B_1 \cup \dots \cup B_n)) > v(B_j \setminus A_i)$, for some i, j . Thus, take $B'_j = (B_j \cap A_i) \cup (A_i \setminus (B_1 \cup \dots \cup B_n))$. Then, $v(B'_j) = v(B_j \cap A_i) + v(A_i \setminus (B_1 \cup \dots \cup B_n)) > v(B_j \cap A_i) + v(B_j \setminus A_i) = v(B_j)$. Additionally, $B'_j \subseteq A_i$, implying by Assumption 1(iii) that (B'_j, A_{-i}) is a feasible allocation. Hence, by Assumption 1(ii), (B'_j, A_{-j}) is also a feasible allocation, and by definition of B'_j , we have as well that $B'_j \cap B_i = \emptyset$, for any $i \neq j$, which implies by Assumption 1(iv) that $(B_1, \dots, B'_j, \dots, B_n) \in \mathcal{F}$, where B_j is replaced by B'_j . But then, by Assumption 2 on \mathcal{F}_G , we must have that $(B_1, \dots, B'_j, \dots, B_n) \in \mathcal{F}_G$. This implies a contradiction because $(B_1, \dots, B_n) \in \mathcal{F}_G^*$, but $\sum_{i=1}^n \gamma_i v(B_i) < \sum_{i=1}^{j-1} \gamma_i v(B_i) + \gamma_j v(B'_j) + \sum_{i=j+1}^n \gamma_i v(B_i)$. \square

Using Lemma A.1, we now prove Theorem 1, which shows an upper bound of $L_\gamma(\mathcal{F}, \mathcal{F}_G)$.

Proof of Theorem 1. Given any \mathcal{F} and \mathcal{F}_G , we know that any allocations $(A_1, \dots, A_n) \in \mathcal{F}^*$ and $(B_1, \dots, B_n) \in \mathcal{F}_G^*$ must satisfy the conditions imposed by Lemma A.1. Let us then consider the following set of scalar variables:

$$\begin{aligned} x_i &= v(A_i \setminus \cup_{k=1}^n B_k), \quad \text{for } i \in N \\ y_i &= v(B_i \setminus \cup_{k=1}^n A_k), \quad \text{for } i \in N \\ w_{ij} &= v(A_i \cap B_j) \quad \text{for } i, j \in N. \end{aligned}$$

Notice that $x_i = v(A_i \setminus \cup_{k=1}^n B_k)$, and that the following equalities hold because of the definition of $v(S)$ for $S \subseteq D$ and the fact that each $A_i \cap A_j = \emptyset$, for every $i, j \in N$:

$$\begin{aligned} y_j + \sum_{k=1}^n w_{kj} - w_{ij} &= v(B_j \setminus \cup_{k=1}^n A_k) + \sum_{\substack{k=1 \\ k \neq i}}^n v(A_i \cap B_k) \\ &= v(B_j \setminus \cup_{k=1}^n A_k) + v\left(\cup_{\substack{k=1 \\ k \neq i}}^n A_i \cap B_k\right) \\ &= v(B_j \setminus A_i). \end{aligned}$$

Hence, we can rewrite the inequalities in (A.1) as

$$x_i \leq \sum_{k=1}^n w_{kj} - w_{ij} + y_j, \quad \text{for } i, j \in N. \quad (\text{A.2})$$

Additionally, because $B_i \cup B_j = \emptyset$ if $i \neq j$, we can express $v(A_i)$ in terms of the scalar variables x_i and w_{ij} , as the following equalities show

$$\begin{aligned} v(A_i) &= v((A_i \setminus \cup_{k=1}^n B_k) \cup (\cup_{k=1}^n A_i \cap B_k)) \\ &= v(A_i \setminus \cup_{k=1}^n B_k) + \sum_{j=1}^n v(A_i \cap B_j) \\ &= x_i + \sum_{j=1}^n w_{ij}. \end{aligned}$$

Equivalently, we can show that $v(B_j) = y_j + \sum_{i=1}^n w_{ij}$. Using these two reformulations of $v(A_i)$ and $v(B_j)$, we can write the expression in (1) that defines $L_\gamma(\mathcal{F}, \mathcal{F}_G)$ as

$$\frac{\sum_{i=1}^n \gamma_i x_i - \sum_{i=1}^n \gamma_i y_i + \sum_{i=1}^n \sum_{j=1}^n w_{ij} (\gamma_i - \gamma_j)}{\sum_{i=1}^n \gamma_i x_i + \sum_{i=1}^n \sum_{j=1}^n w_{ij} \gamma_j}. \tag{A.3}$$

Notice that written in this way it is clear that $L_\gamma(\mathcal{F}, \mathcal{F}_G)$ is a Fractional Linear function of variables x , y , and w . Because for any \mathcal{F} and \mathcal{F}_G , the inequalities of (A.2) must hold, then we can find an upper bound on $L_\gamma(\mathcal{F}, \mathcal{F}_G)$ for any \mathcal{F} and \mathcal{F}_G , given γ , by solving the following Fractional Linear Program:

$$\begin{aligned} \text{Maximize} \quad & \frac{\sum_{i=1}^n \gamma_i x_i - \sum_{i=1}^n \gamma_i y_i + \sum_{i=1}^n \sum_{j=1}^n w_{ij} (\gamma_i - \gamma_j)}{\sum_{i=1}^n \gamma_i x_i + \sum_{i=1}^n \sum_{j=1}^n w_{ij} \gamma_i} \\ \text{subj. to} \quad & x_i + w_{ij} - \sum_{k=1}^n w_{kj} - y_j \leq 0 \quad \text{for } i, j \in N \tag{A.4} \\ & x_i \geq 0 \quad \text{for } i \in N \\ & y_i \geq 0 \quad \text{for } i \in N \\ & w_{ij} \geq 0 \quad \text{for } i, j \in N. \end{aligned}$$

Given that the constraints in the maximization problem are all homogeneous, we can rewrite this problem, by scaling all the variables, into the following equivalent linear program:

$$\begin{aligned} \text{Maximize} \quad & \sum_{i=1}^n \gamma_i x_i - \sum_{i=1}^n \gamma_i y_i + \sum_{i=1}^n \sum_{j=1}^n w_{ij} (\gamma_i - \gamma_j) \\ \text{subj. to} \quad & \sum_{i=1}^n \gamma_i x_i + \sum_{i=1}^n \sum_{j=1}^n w_{ij} \gamma_i = 1 \\ & x_i + w_{ij} - \sum_{k=1}^n w_{kj} - y_j \leq 0 \quad \text{for } i, j \in N \tag{A.5} \\ & x_i \geq 0 \quad \text{for } i \in N \\ & y_i \geq 0 \quad \text{for } i \in N \\ & w_{ij} \geq 0 \quad \text{for } i, j \in N. \end{aligned}$$

The optimization problem in (A.5) is a linear program with variables x , y , and w and an objective function bounded above by one. Thus, we know that strong duality must hold,

and we can find the desired upper bound by studying the dual linear program:

$$\begin{aligned} \text{Minimize} \quad & s \\ \text{subj. to} \quad & s\gamma_i + \sum_{j=1}^n \lambda_{ij} \geq \gamma_i \quad \text{for } i \in N \\ & \gamma_j - \gamma_k + s\gamma_k - \sum_{\substack{i=1 \\ i \neq k}}^n \lambda_{ij} \geq 0 \quad \text{for } k, j \in N \tag{A.6} \\ & \gamma_j \geq \sum_{i=1}^n \lambda_{ij} \quad \text{for } j \in N \\ & \lambda_{ij} \geq 0 \quad \text{for } i, j \in N. \end{aligned}$$

Problem (A.6) in turn can be rewritten as

$$\begin{aligned} \text{Minimize} \quad & \max \left\{ \left\{ 1 - \frac{1}{\gamma_i} \sum_{j=1}^n \lambda_{ij} \right\}_{i=1}^n, \left\{ \frac{1}{\gamma_k} \left(\sum_{\substack{i=1 \\ i \neq k}}^n \lambda_{ij} + \gamma_k - \gamma_j \right) \right\}_{k,j=1}^n \right\} \\ \text{subj. to} \quad & \gamma_j \geq \sum_{i=1}^n \lambda_{ij} \quad \text{for } j \in N \\ & \lambda_{ij} \geq 0 \quad \text{for } i, j \in N. \tag{A.7} \end{aligned}$$

Let us define $f(\gamma) : [\gamma_{\min}, \gamma_{\max}]^n \rightarrow [0, 1]$ as the optimal value of the dual problem (A.7). Now we can see that each term that appears in the objective of problem (A.7) is a quasi-convex function of both γ and λ ; thus, the objective itself is quasiconvex in these variables because it is the finite maximization of quasiconvex functions. Moreover, the feasible region of problem (A.7) is convex in γ and λ , which means that the function $f(\gamma)$ must be quasiconvex in γ because it is the minimization of a quasiconvex function over a convex set. But then, if we wish to find $\max_{\gamma \in [\gamma_{\min}, \gamma_{\max}]^n} f(\gamma)$, then we need only to look at the extreme points of the hypercube $[\gamma_{\min}, \gamma_{\max}]^n$ (see Bertsekas et al. 2003 for a proof of this result).

Now, as we prove in Lemma A.2, if we take the instance where the first n_0 values of γ are γ_{\max} , and the rest are γ_{\min} , for $n_0 \in N$, we can see that the optimal value is $\max\{\delta, \frac{n-1}{n+n_0+(1-\delta)(n-n_0)-1}\}$. But, notice that this is a decreasing function of n_0 , implying that the instance of γ that maximizes the solution to problem (A.4) is when $n_0 = 1$. In this case, we recover $\max\{\delta, \frac{n-1}{n+(1-\delta)(n-1)}\}$, which proves the theorem. \square

Lemma A.2. *The optimal value of the linear program (A.5), when $\gamma_i = \gamma_{\max}$ for all $i \in \{1, \dots, n_0\}$, and $\gamma_i = \gamma_{\min}$ for all $i \in \{n_0 + 1, \dots, n\}$, for all $n_0 \in \{1, \dots, n\}$ is*

$$\max \left\{ \delta, \frac{n-1}{n+n_0+(1-\delta)(n-n_0)-1} \right\},$$

where $\delta = \frac{\gamma_{\max} - \gamma_{\min}}{\gamma_{\max}}$.

Proof. Notice that the linear program (A.6) is a dual of the linear program (A.5). Thus, we will produce a primal and a dual feasible instance, both attaining the proposed optimal value, which will show that it is indeed the optimal value. For ease of notation, we will define $X = \frac{n-1}{n+n_0+(1-\delta)(n-n_0)-1}$.

For this, we consider first the case where $\delta > X$.

In this case, consider the following primal feasible point:

$$x_i = y_i = 0, \forall i \in N, \quad w_{1,n_0+1} = \frac{1}{\gamma_{\max}},$$

$$w_{ij} = 0, \forall (i, j) \neq (1, n_0 + 1).$$

By simply replacing these values in problem (A.5), we can see that the objective of δ is achieved.

For the dual problem, let us consider the following feasible point:

$$s = \delta, \quad \lambda_{ij} = 0 \text{ for } j \in \{(n_0 + 1), \dots, n\}$$

$$\lambda_{ij} = \frac{\gamma_{\min}}{n_0} \text{ for } i, j \in \{1, \dots, n_0\}, \quad \lambda_{ij} = \frac{\gamma_{\min}^2}{\gamma_{\max} n_0}$$

$$\text{for } i \in \{(n_0 + 1), \dots, n\} \text{ and } j \in \{1, \dots, n_0\}.$$

By evaluating the dual problem in this specific point, we can see that it achieves an objective value of δ , and it is feasible only when $\delta \geq X$.

In the case when $X \geq \delta$, we take the following primal feasible solution:

$$x_1 = 0, \quad x_i = \frac{X}{(n-1)\gamma_{\max}}, \quad \forall i \neq 1, \quad y_i = 0, \forall i \in N$$

$$w_{1j} = \frac{X}{(n-1)\gamma_{\max}}, \forall j \in N, \quad w_{ij} = 0, \quad \forall i \in \{2, \dots, n\}, \quad j \in N.$$

Simple algebra will show that this solution is primal feasible and achieves the objective value of X .

Finally, we take the following dual solution:

$$s = X, \quad \lambda_{ij} = \frac{\gamma_{\max}(X - \delta)}{(n-1)}, \forall j \in \{(n_0 + 1), \dots, n\}, \quad i \in N.$$

$$\lambda_{ij} = \frac{\gamma_{\max}}{n_0} \left(1 - \frac{(n-n_0)(X - \delta)}{(n-1)} - X \right), \forall i, j \in \{1, \dots, n_0\}$$

$$\lambda_{ij} = \frac{1}{n_0} \left(\gamma_{\min}(1 - X) - \frac{\gamma_{\max}(n-n_0)(X - \delta)}{(n-1)} \right),$$

$$\forall i \in \{(n_0 + 1), \dots, n\}, \quad j \in \{1, \dots, n_0\}.$$

With some algebra, this solution can be seen to be dual feasible when $X \geq \delta$, and it clearly achieves an objective values of X because $s = X$.

Hence, both when $X \geq \delta$ and when the converse occurs, we have produced dual and primal feasible solutions that achieve the objective value of $\max\{X, \delta\}$, proving that this must indeed be the optimal value. \square

Now that we have proved Theorem 1, we proceed to prove Theorem 2. For this, we need only to show that we can asymptotically approximate the upper bound proven in Theorem 1.

Proof of Theorem 2. We need only to prove that there exists instances of \mathcal{F} , \mathcal{F}_G , and γ such that $L_\gamma(\mathcal{F}, \mathcal{F}_G)$ achieves the values in (6) and (7). For this, we will generalize Instances 1 and 2, for n agents.

We begin by considering an instance family that achieves $L_\gamma(\mathcal{F}, \mathcal{F}_G) = \delta$, for any n and δ .

Instance A.1. Take $D = \{d_1\}$, such that $v(d_1) = 1$, any n , and $\gamma_1 = \gamma_{\max}$, $\gamma_i = \gamma_{\min} \leq \gamma_{\max}$ for all $i \in \{2, \dots, n\}$. Given this D with only one job, consider $\mathcal{F} = \{\mathbf{A}, \mathbf{B}\}$, where $A_1 = \{d_1\}$, $A_i = \emptyset$ for all $i \in \{2, \dots, n\}$, and $B_1 = \emptyset$, $B_2 = \{d_1\}$, and $B_i = \emptyset$ for all $i \in \{3, \dots, n\}$. Finally, if we take $\mathcal{F}_G = \{\mathbf{B}\}$, then the only efficient allocations would be \mathbf{A} that gives a value of γ_{\max} , whereas by definition, the only allocation in \mathcal{F}_G would be \mathbf{B} , which implies that $L_\gamma(\mathcal{F}, \mathcal{F}_G) = \frac{\gamma_{\max} - \gamma_{\min}}{\gamma_{\max}} = \delta$.

Now we present a family of instances that have $L_\gamma(\mathcal{F}, \mathcal{F}_G) = \frac{n-1}{n+(1-\delta)(n-1)} - \epsilon$, for any $\epsilon > 0$, where $\mathcal{F}_G = \mathcal{F}_G^{\text{mM}}$ is the restriction to only max-min fair allocations. In Instance 2, we presented an instance for two agents given by three jobs that had the following properties: two of the jobs could be fulfilled by any single provider, whereas one of the jobs overlapped with all the other jobs and thus could only be assigned by itself to a provider. We will generalize these properties now to $2n - 1$ jobs.

Instance A.2. Given any n , let $D = \{d_1, \dots, d_n, \dots, d_{2n-1}\}$, such that $v(d_i) = 1$, for all $i \in \{1, \dots, n\}$, and $v(d_j) = 1 - \kappa$, for $j \in \{n+1, \dots, 2n-1\}$. Let $\gamma_1 = \gamma_{\max}$ and $\gamma_i = \gamma_{\min}$, for all $i \in \{2, \dots, n\}$. A subset $A \subseteq D$ can be assigned to a single provider if either $d_j \notin A$, for all $j \in \{n+1, \dots, 2n-1\}$, or $A = \{d_i\}$ for some $i \in \{n+1, \dots, 2n-1\}$. Let \mathcal{F} be formed by all possible disjoint combinations of such subsets of D . Let $p: \mathcal{P}(D) \rightarrow \mathbb{R}$ be a strict monotonic payment function satisfying (3) and

$$v(S) > v(T) \Rightarrow p(S) > p(T), \forall S, T \in \mathcal{P}(D).$$

Let $\mathcal{F}_G^{\text{mM}}(p, N)$ be the associated max-min fair restriction, as described in (4).

Given this instance, the only efficient allocation would be \mathbf{A} with $A_1 = \{d_1, \dots, d_n\}$, $A_i = \{d_{n+i-1}\}$ for $2 \leq i \leq n$. That is, we assign all of the first n jobs to one provider and distribute the remaining jobs between the remaining providers. Moreover, the only allocation in $\mathcal{F}_G^{\text{mM}}$ (modulo symmetries) is given by \mathbf{B} with $B_i = \{d_i\}$, for $1 \leq i \leq n$. As in the $n = 2$ case, this is because of smaller $1 - \kappa$ value generated by the last $n - 1$ jobs and the monotonicity of the payment function $p(\cdot)$. Hence, this instance generates a loss $L_\gamma(\mathcal{F}, \mathcal{F}_G) = \frac{(n\gamma_{\max} + \gamma_{\min}(n-1)(1-\kappa)) - (\gamma_{\max} + (n-1)\gamma_{\min})}{n\gamma_{\max} + \gamma_{\min}(n-1)(1-\kappa)} = \frac{(n-1)\gamma_{\max} - \kappa\gamma_{\min}(n-1)}{n\gamma_{\max} + \gamma_{\min}(n-1)(1-\kappa)}$
 $\xrightarrow{\kappa \rightarrow 0} \frac{n-1}{n+(1-\delta)(n-1)}$. Therefore, for any $\epsilon > 0$, there exists a κ small enough such that $L_\gamma(\mathcal{F}, \mathcal{F}_G) = \frac{n-1}{n+(1-\delta)(n-1)} - \epsilon$.

Instance A.3. Given any n , and t , positive integers, let $D = \cup_{k=1}^t C^k \cup S$, where $C^k = \{d_2^k, \dots, d_n^k\}$, and $S = \{d_1^s, \dots, d_{n(t+1)}^s\}$. Let as well $\gamma_1 = \gamma_{\max} \geq \gamma_{\min} = \gamma_i$, for all $i \in \{2, \dots, n\}$. A subset of jobs $A \subseteq D$ can be assigned to a single provider if either $A \cap \cup_{k=1}^t C^k = \emptyset$ or $|(A \cap C^k)| \leq 1$, for each $k \in \{1, \dots, t\}$ and $A \cap S = \emptyset$. Let \mathcal{F} be formed by all possible disjoint combinations of such subsets of D . Let, as well $v(d) = 1$, for all $d \in D$, and $p: \mathcal{P}(D) \rightarrow \mathbb{R}$ be a strict monotonic payment function satisfying both (3) and

$$v(S) > v(T) \Rightarrow p(S) > p(T), \forall S, T \in \mathcal{P}(D).$$

Let $\mathcal{F}_G^{\text{mM}}(p, N)$ be the associated max-min fair restriction, as described in (4).

Given this instance, the only efficient allocation (modulo symmetries) would be \mathbf{A} , such that $A_1 = S$, $A_i = \{d_i^1, d_i^2, \dots, d_i^t\}$, for all $i \in \{2, \dots, n\}$. That is, we assign all the jobs in the set S to the provider that generates the highest value, and we assign one job of each C^k to the rest of the providers, for a total of t jobs. Moreover, the only allocations in $\mathcal{F}_G^{\text{mm}}$ are of the form \mathbf{B} , such that $B_i \subseteq S$, and $|B_i| = t + 1$, for each $i \in N$. In other words, we divide the $(t + 1)n$ jobs of S among all the providers equally. In the efficient allocation all providers, except for the first one, are being allocated exactly t jobs, whereas in any max-min fair allocation, all providers are being allocated exactly $t + 1$ jobs. Hence, the loss generated by this instance, which comes from the fact that none of the jobs in $\cup_{k=1}^t C^k$ are allocated for any allocation in $\mathcal{F}_G^{\text{mm}}$, is $\frac{(t+1)n\gamma_{\max} + t(n-1)\gamma_{\min} - (t+1)(\gamma_{\max} + (n-1)\gamma_{\min})}{(t+1)n\gamma_{\max} + t(n-1)\gamma_{\min}} = \frac{t(n-1) + \delta(n-1)}{t(n+(n-1)(1-\delta)) + n}$. Therefore, for any $\epsilon > 0$, there exists a t large enough such that $L_\gamma(\mathcal{F}, \mathcal{F}_G) = \frac{n-1}{n+(1-\delta)(n-1)} - \epsilon$.

To better visualize the structure of Instances A.2 and A.3, we can imagine a graph on the elements of D , where a set of jobs can be assigned together only when there is no edge between any pair of the corresponding vertices. Thus, in the case of Instance A.2, there would be no edges between any pair of the first n vertices, whereas every pair of the last $n - 1$ vertices would be joined by an edge. Finally, every vertex of the last $n - 1$ will be adjacent to all of the first n vertices. An example for three providers of this graph can be seen in Figure A.1. These types of graphs are known as complete split graphs (see Le and Peng 2015). Similarly, in the case of Instance A.3, there would be no edges between any pair of vertices in the set S , whereas each pair of vertices in the same C^k would be connected by an edge. Moreover, every vertex in S would be connected to every vertex in each of the C^k . An example for three providers and $t = 1$ can be seen in Figure A.1. An allocation of the jobs can be seen as a covering of this graph by independent sets (sets of vertices without any edge joining two vertices of the set).

The family of Instances A.1–A.3 proves that $L_\gamma(\mathcal{F}, \mathcal{F}_G)$ can be taken as close to $\max\{\delta, \frac{n-1}{n+(1-\delta)(n-1)}\}$ as desired. This concludes the proof of Theorem 2. \square

We now present Instance A.4, which shows how when we relax Assumption 1(ii), we can achieve a loss that grows asymptotically to 100% with the number of providers.

Instance A.4. Consider $n = 2k - 1$, for any integer $k > 0$, $D = \{d_1^1, \dots, d_k^1, \dots, d_1^k, \dots, d_k^k\}$, and

$$\mathcal{F} = \{A \mid A_i \subseteq \{d_1^i, \dots, d_k^i\}, \text{ for } i \in \{1, \dots, k\}, \\ A_j \subseteq \{d_{j-k}^1, \dots, d_{j-k}^k\}, \text{ for } j \in \{k + 1, \dots, 2k - 1\}\}.$$

Let $v(d) = \frac{1}{k}$, for $d \in D$, let $\gamma_i = 1$ for $i \in \{1, \dots, k\}$, and $\gamma_j = \frac{1}{k}$ for $j \in \{k + 1, \dots, 2k - 1\}$, and let $p_i(A_i) = \gamma_i \sum_{d \in A_i} v(d)$. Then,

$$L_\gamma(\mathcal{F}, \mathcal{F}_G^{\text{mm}}(p, N)) = 1 - \frac{4n}{(n + 1)^2},$$

where $\mathcal{F}_G^{\text{mm}}$ is the allocations that satisfy max-min fairness, as defined in (4).

To see this, notice that the only max-min fair allocation, that would leave each provider with a $p(B_i) = \frac{1}{k}$, for $i \in N$, is (modulo permutations) \mathbf{B} , such that $B_i = \{d_k^i\}$, for $i \in \{1, \dots, k\}$, and $B_j = \{d_{j-k}^1, \dots, d_{j-k}^k\}$, for $j \in \{j + 1, \dots, 2k - 1\}$. This allocation would lead to a total value of $\frac{2k-1}{k}$. On the other hand, the value-maximizing allocation would only allocate jobs to the first k providers, by taking \mathbf{A} such that $A_i = \{d_1^i, \dots, d_k^i\}$, for $i \in \{1, \dots, k\}$. This would lead to a total value generated of exactly k , which would imply that

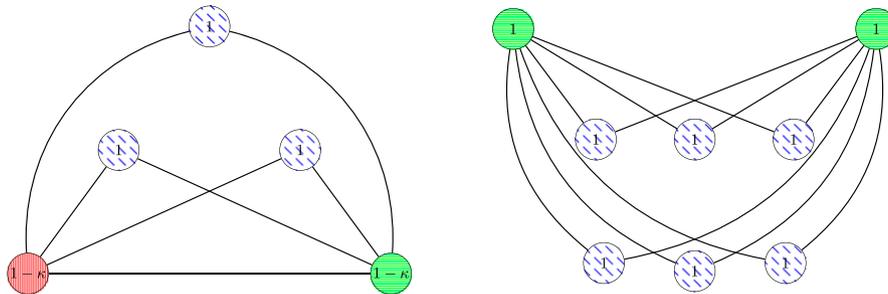
$$L_\gamma(\mathcal{F}, \mathcal{F}_G^{\text{mm}}(p, N)) = \frac{k - \frac{2k-1}{k}}{k} \\ = 1 - \frac{4n}{(n + 1)^2}.$$

Notice finally that this instance does not satisfy Assumption 1(ii) because the first k providers can complete any subset of trips from $\{d_1^i, \dots, d_k^i\}$, for provider $i \in \{1, \dots, k\}$, but the last $k - 1$ providers can only perform subsets of $\{d_{j-k}^1, \dots, d_{j-k}^k\}$, for provider j in $\{k, \dots, 2k - 1\}$. Therefore, almost any permutation of a feasible allocation would lead to an infeasible allocation.

We will now prove Propositions 3–7 from Section 4.

Proof of Proposition 7. In order to prove this proposition, we will first formally define the set of guarantees \mathcal{F}_G^c , as any subset of \mathcal{F}^c , that satisfies Assumption 2, replacing in the definition of the assumption \mathcal{F} by \mathcal{F}^c , and using the generalized notion of $v(C_i)$ for $C \in \mathcal{F}^c$, described in Section 4.

Figure A.1. (Color online) Feasibility Graphs for Instance A.2 (Left Panel), When n Is Three, and Instance A.3 (Right Panel), When n Is Three and t Is One



In particular, we extend the notion of uniform income guarantees under monotonic payment functions: given $p(\cdot)$, a monotonic payment function, we take $\mathcal{F}_G^c = \{C \in \mathcal{F}^c \mid p(C_i) \geq \tau, \text{ for } i \in N\}$. As in Section 2.4, it is easy to see that these guarantees satisfy the extended version of Assumption 2. Now, we will prove that given any uniform income guarantee, $L_\gamma(\mathcal{F}^c, \mathcal{F}_G^c) = 0$.

Consider any allocation $A \in \mathcal{F}^*$. We know that any permutation A_σ of A is as well in \mathcal{F} ; hence, take $C \in \mathcal{F}^c$, such that $C = (\{\theta_\sigma\}_{\sigma \in S_n}, \{A_\sigma\}_{\sigma \in S_n})$, where S_n is the symmetric group of all permutations of N , $\theta_\sigma = \frac{1}{n!}$, and A_σ is a specific permutation of A . Hence, $v(C_i) = v(C_j) = \sum_{i=1}^n \frac{1}{n} v(A_i)$, for each $i \neq j \in N$. This implies that C must be in any nonempty uniform income guarantee. To see this, let us assume by contradiction that there is a nonempty \mathcal{F}_G^c such that $C \notin \mathcal{F}_G^c$. Without loss of generality, because they both induce the same ordering on the subsets of D , we will assume that $p(\cdot) = v(\cdot)$. Hence, $C \notin \mathcal{F}_G^c$ implies that the corresponding income guarantee, τ is greater than $\sum_{i=1}^n \frac{1}{n} v(A_i)$. But, then there must exist at least one $B \in \mathcal{F}^c$ such that $v(B_i) \geq \tau > \sum_{i=1}^n \frac{1}{n} v(A_i)$, for each B_i , which implies that $\sum_{i=1}^n v(A_i) < \sum_{i=1}^n v(B_i) = \sum_{i=1}^n \sum_{j=1}^k \theta_j v(B_i^j) = \sum_{j=1}^k \theta_j \sum_{i=1}^n v(B_i^j) \leq \max_{j=1}^k \sum_{i=1}^n v(B_i^j)$. But then, there exists an allocation $B^j \in \mathcal{F}$ that achieves a higher total value than $A \in \mathcal{F}^*$, which leads to a contradiction and proves the proposition. \square

Proof of Proposition 3. Let $\hat{k} = \sum_{i=1}^n k_i$, and let us assume that $D = \{d_1, \dots, d_m\}$, where jobs are ordered decreasingly in $v(d_i)$; then, for any allocation $A \in \mathcal{F}^*$, $\sum_{i=1}^n v(A_i) = \sum_{j=1}^{\hat{k}} v(d_j)$.

To see this, notice that the total number of jobs that can be allocated is \hat{k} , and thus, if the total value generated in any $A \in \mathcal{F}^*$ was less than $\gamma_{\max} \sum_{j=1}^{\hat{k}} v(d_j)$, then take the job with smallest intrinsic value being allocated, and replace it by the job with highest intrinsic value in $\{d_1, \dots, d_{\hat{k}}\} \setminus (\cup_{i=1}^n A_i)$. This replacement would generate a feasible allocation, and would improve the total value generated, which leads to a contradiction because $A \in \mathcal{F}^*$.

Now to prove that $L_\gamma(\mathcal{F}, \mathcal{F}_G) = 0$, we will proceed in two steps. First, we will show that any allocation $B \in \mathcal{F}$ can be Pareto dominated, in the sense of Assumption 2, by an allocation A in \mathcal{F} , that uses only jobs in $\{d_1, \dots, d_{\hat{k}}\}$; the second is that any allocation $C \in A$ that uses only jobs in $\{d_1, \dots, d_{\hat{k}}\}$ can be Pareto dominated by an allocation C in \mathcal{F} , that uses all jobs in $\{d_1, \dots, d_{\hat{k}}\}$. By transitivity of the Pareto dominance, this will imply that any allocation can be Pareto dominated by an allocation that uses all elements in $\{d_1, \dots, d_{\hat{k}}\}$, and therefore is in \mathcal{F}^* , which, by Assumption 2, will imply that there is an element of \mathcal{F}^* in \mathcal{F}_G ; therefore, $L_\gamma(\mathcal{F}, \mathcal{F}_G) = 0$.

Take any allocation $B \in A$, if $(\cup_{i=1}^n B_i) \setminus \{d_1, \dots, d_{\hat{k}}\} = \emptyset$, then B allocates only elements of $\{d_1, \dots, d_{\hat{k}}\}$. Otherwise, consider A such that we replace in B every job in $(\cup_{i=1}^n B_i) \setminus \{d_1, \dots, d_{\hat{k}}\}$ by an element in $\{d_1, \dots, d_{\hat{k}}\} \setminus (\cup_{i=1}^n A_i)$. Because every job we replaced must have a lower intrinsic value than any job in the first \hat{k} , then we know that $v(B_i) \leq v(A_i)$, for each $i \in N$.

Now, assume we have a $A \in \mathcal{F}$, such that only jobs in $\{d_1, \dots, d_{\hat{k}}\}$ are allocated; then if there are any jobs in the first \hat{k} not allocated in C , this means that there is at least one provider i such that $|A_i| < k_i$. Consider then the allocation

$C \in \mathcal{F}$, such that we add jobs from $\{d_1, \dots, d_{\hat{k}}\}$ to A , until all $|A_i| = k_i$. This allocation C Pareto dominates allocation A , and uses exactly all elements in $\{d_1, \dots, d_{\hat{k}}\}$. Hence, as mentioned, this proves that $L_\gamma(\mathcal{F}, \mathcal{F}_G) = 0$, for any \mathcal{F}_G , satisfying Assumption 2. \square

Proof of Proposition 4. We begin by proving (i). Without loss of generality, we will normalize $v(d) = 1$, for all $d \in D$. To show that we have zero loss under any uniform income guarantee \mathcal{F}_G , we will show that when $|D| < n$, $\max_{A \in \mathcal{F}} \min_{i=1}^n v(A_i) = 0$, and when $|D| \geq n$, $\max_{A \in \mathcal{F}} \min_{i=1}^n v(A_i) = 1$.

When $|D| < n$, any allocation $A \in \mathcal{F}$ will necessarily have an $A_i = \emptyset$, which implies that $\max_{A \in \mathcal{F}} \min_{i=1}^n v(A_i) = 0$. On the other hand, if $|D| \geq n$, then for any $A \in \mathcal{F}$, each A_i is either $\{d_j^c\}$, for some $d_j^c \in C$, or $A_i \subseteq S$. In the first case, $v(A_i) = 1$; in the second case, $v(A_i) = |A_i|$. Thus, the only way of having $\min_{i=1}^n v(A_i) > 1$ would be if each $A_i \subseteq S$, and $|A_i| \geq 2$, for all $i \in N$, but this implies that $|S| \geq 2n$, which in turn, contradicts our hypothesis. Therefore, $\max_{A \in \mathcal{F}} \min_{i=1}^n v(A_i) = 1$.

In both cases outlined, \mathcal{F}^* will always contain an allocation A , such that $\min_{i=1}^n v(A_i) = \max_{B \in \mathcal{F}} \min_{i=1}^n v(B_i)$, which implies that $L_\gamma(\mathcal{F}, \mathcal{F}_G) = 0$, for any uniform income guarantee under monotonic payment functions, \mathcal{F}_G .

To show (ii), we simply observe that Instance 5 is a generalized form of Instance 2. In particular, by taking $v(d) = 1 - \kappa$, for all $d \in C$, $v(d) = 1$, for all $d \in S$, and taking $|C| = n - 1$, and $|S| = n$, we obtain exactly Instance 2, when $n = 2$, and Instance A.2 when $n \geq 2$. \square

We now describe Example A.1, which shows how the dependency of the loss on the variance of the values described in Example 1 extends to the n provider case.

Example A.1. Consider a variant of Instance A.2, where the κ term is taken to be a random variable, $\kappa \sim U[-\frac{\Delta}{2}, \frac{\Delta}{2}]$, and $\gamma_{\min} = \gamma_{\max} = 1$. Thus, the value $v_i = v(d_i) = 1 - \kappa \sim U[1 - \frac{\Delta}{2}, 1 + \frac{\Delta}{2}]$, for each $i \in \{n + 1, \dots, 2n - 1\}$. Hence, if we take the expectation of the loss, with respect to the error κ , we get

$$\begin{aligned} \mathbb{E}_\kappa(L_\gamma(\mathcal{F}, \mathcal{F}_G)) &= \int_0^{\min\{1, \frac{\Delta}{2}\}} \frac{1}{\Delta} \frac{(n-1)(1-\kappa)}{n+(n-1)(1-\kappa)} d\kappa \\ &= \frac{1}{\Delta} \left(\frac{n}{n-1} \log \left((n-1) \min \left\{ 1, \frac{\Delta}{2} \right\} - (2n-1) \right) \right. \\ &\quad \left. + \min \left\{ 1, \frac{\Delta}{2} \right\} - \frac{n}{n-1} \log(2n-1) \right) \\ &:= g(\Delta). \end{aligned}$$

As in Example 1, it can be seen that $g(\Delta)$ is decreasing in Δ , which implies that it is decreasing in the variance of the v_i values: for higher variance, there is a lower expected loss.

Proof of Proposition 5. We will first prove that if $n = 1$ and $\gamma_{\min} = \gamma_{\max}$, then $L_\gamma(\mathcal{F}, \mathcal{F}_G) = 0$, for any set of allocations with guarantees, \mathcal{F}_G . To show this, we simply observe that because of Assumption 2, any $A \in \mathcal{F}^*$ must also satisfy $A \in \mathcal{F}_G$ because clearly for any $B \in \mathcal{F}_G$, $v(B_1) \leq v(A_1)$.

Now, we will show that the loss is zero when $n \geq |D|$, for any set of allocations with guarantees. For this, we assume without loss of generality that it is always feasible to allocate at least one job to any specific provider (if not, then

there is a job that cannot be completed by any provider, and we could then simply ignore it). Hence, we claim that $\max_{A \in \mathcal{F}} \sum_i v(A_i) = \sum_{d \in D} v(d)$. This is because when $n \geq |D|$, we can always allocate all jobs by allocating one job per provider to the first $|D|$ providers.

Now, we claim that in \mathcal{F}_G^* there exists an allocation B , such that $D \subseteq \cup_i B_i$. To see why this is, assume to the contrary that no such allocation exists. Then, take any $A \in \mathcal{F}_G^*$, there exists thus $d \in D$ such that $d \notin \cup_i A_i$. Moreover, because $|D| \leq n$, then there exists a provider i , such that $A_i = \emptyset$. Hence, simply take A' such that $A'_j = A_j$, for $j \neq i$, and $A'_i = \{d\}$. This leads to a contradiction, because by Assumption 2, $A' \in \mathcal{F}_G$, and $\sum_{i \in N} v(A'_i) > \sum_{i \in N} v(A_i)$, but $A \in \mathcal{F}_G^*$. Therefore, there exists an allocation $B \in \mathcal{F}_G^*$, such that $D \subseteq \cup_i B_i$, and thus, $\sum_{i \in N} v(B_i) = \sum_{d \in D} v(d)$, which implies that $L_\gamma(\mathcal{F}, \mathcal{F}_G) = 0$, for any set of allocations with guarantees \mathcal{F}_G . \square

Proof of Proposition 6. Without loss of generality, we can assume that all intrinsic values are 1: that is, $v(d) = 1$, for each $d \in D$. Now, because $|D| < 2n$, and because we can always allocate only one job to any provider, then $\max_{A \in \mathcal{F}} \min_{i \in N} v(A_i) = 1$. Hence, any uniform income guarantee can at most guarantee the payment produced by exactly one job. Therefore, because there always exists an allocation in \mathcal{F}^* that allocates at least one job to each provider, we conclude that $L_\gamma(\mathcal{F}, \mathcal{F}_G) = 0$, under any uniform income guarantee under monotonic payment functions. \square

Appendix B. Numerical Analysis of Synthetic and Real-World Data

In this section, we provide the details of the numerical analysis we discuss in Section 5. The objective of this analysis is to demonstrate both the magnitude of the relative loss and the different drivers of this loss in a particular setting covered by our general theoretical analysis.

B.1. Instances Generated with Real-World Data

We used the publicly available data set provided by NYC Taxi and Limousine Commission (2016). This data set

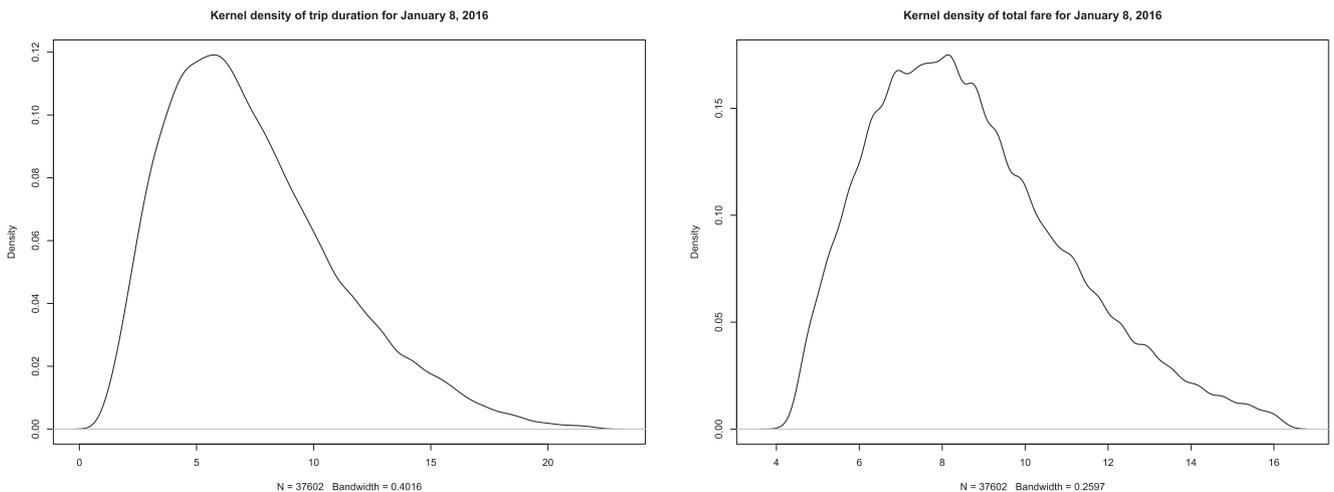
includes, for each yellow taxi ride, the total fare, the starting and ending locations as well as total time of the ride. We considered several dates (from January 4 to January 8, 2016). For each of these dates, we looked at the trips that started between 9 a.m. and 5 p.m., in order to restrict our attention to a time horizon with a relatively constant rate of trips per time.

Moreover, in order to better conform to part (ii) of Assumption 1, we took two different approaches. First, we filtered the trips that started and ended in a limited region of Manhattan (we took Midtown, Upper West Side, and Upper East Side). By considering this small geographical region, we limit the effect of spacial considerations, and better conform to part (ii) of Assumption 1, that any feasible set of jobs could be performed by any provider. For our second approach, although we did not limit the starting and ending region, we added geographical constraints to our allocations that we describe. By adding these spacial considerations, we ensure that each set of feasible trips is geographically consistent, while at the same time, we satisfy part (ii) of Assumption 1 by allowing any provider to perform any set of feasible trips.

In Figure B.1, we can see the empirical distribution of the total trip duration and total fare paid, for Midtown Manhattan, on January 8, 2016. In particular, for this specific date we can see that the mean in total duration for this region is of 7.41, whereas the variance is 13.48. At the same time the mean of the total fares is \$8.68, and the variance is \$5.63. Finally, we also cleaned the data by removing the trips in the top 0.1% of both total time elapsed and total fare, thus removing several outliers that were clearly because of corrupted data (trips of almost 24 hours or more than \$1,000).

In order to compute the average relative loss when considering only max-min fair solutions, we generated demand instances using these data. We partitioned the time horizon into intervals of w minutes, and from each of these intervals, we sampled 30 trips uniformly at random. We considered different values of w , from 10 to 20 minutes. We defined feasible allocations in two ways. For our first approach, we only required that trips that intersected in time were not allocated to the same provider. For our second

Figure B.1. Kernel Density of Trip Duration for January 8, 2016



Notes. (Left panel) Empirical distribution of the duration of trips. (Right panel) Empirical distribution of the total fare of trips.

approach, we added the restriction that two trips can only be allocated together if, driving at an average speed of v mph, a driver can reach the start point of one trip from the end point of the other trip. We considered speeds of v in the range $\{7.44\text{mph}, 10.7\text{mph}\}$, which were the mean and median speeds across the whole city in 2016 (see NYC Department of Transportation 2018). Then, we solved for both the total value-maximizing solution and for the value-maximizing solution among the max-min fair solutions, for a varying number of providers. For this, we used an Integer Programming formulation of the allocation problems. In order to obtain the max-min fair allocations, we first solved for the max-min objective and then constrained the allocations to ensure that all providers received at least that amount of total fare. We limited ourselves to 30 jobs because of computational considerations (solving for the max-min solution is NP hard in general). After we obtained the two value-maximizing solutions, we computed the total relative loss across the whole time horizon by taking the relative difference of total value with and without the max-min fair restriction. We sampled the instances 100 times and computed for each number of providers the average value loss across these 100 samples.

In Figures B.2 and B.3, we can see the relative value loss as a function of the number of providers, without any geographical restrictions, for different combinations of dates, regions, and size of interval, w . In Figure B.4, we can see the relative value loss as a function of the number of providers, under the feasibility constraints mentioned, for the whole city of New York, for different average speeds v . We show here a representative set of our results; the complete set of results is available upon request.

We can see in Figure B.2 that there does not seem to be much difference from region to region. We observed similar results across the three regions for all the combinations of dates and values of w we tested. In Figure B.3, we see that for the same region but different dates the relative losses do not seem to change much. Nevertheless, by comparing Figure B.2 with Figure B.3, we observe that the losses do seem to increase when w is decreased. This is consistent with the fact that we are taking the same number of jobs in both, resulting in a higher density of trips per time when we decrease w .

In both Figures B.2 and B.3, we observed the same pattern mentioned in Section 5 regarding the effect of the providers to job ratio on the loss, namely for extreme values of this ratio the loss collapses to zero, and the loss achieves its maximum value at an intermediate value. Moreover, for every day and region we analyzed, we observe that the curves of maximum losses and average losses are relatively close together, implying that the low average losses are because of frequent low losses, as opposed to infrequent high losses.

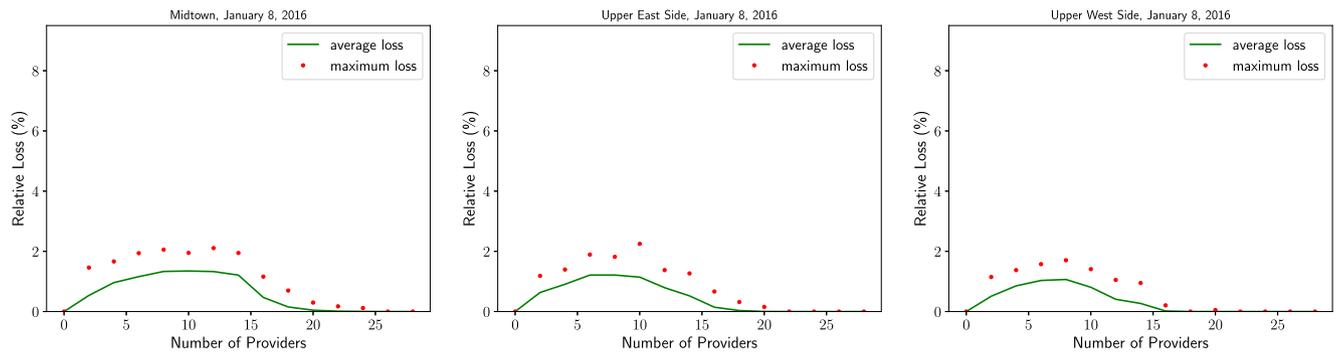
In Figure B.4, we can see that adding a geographic feasibility constraint does not seem to change significantly the average losses. Furthermore, changing the average speed from the mean bus speed in the city to the median bus speed does not result in any apparent change in the average losses. It is worth noting that increasing the speed v did lead eventually to almost the same results seen in

Figure B.3, whereas reducing v to zero led to a constant zero loss.

As a robustness test of how our results would change when Assumption 1(ii) is relaxed, we also computed the average loss in instances that violate the symmetry assumption. Namely, we parametrized the symmetry of each instance by s : in an instance with symmetry $s\%$, each provider can only perform $s\%$ of all the jobs. To construct instances with $s\%$ symmetry, we took the same set of jobs we considered in our first approach defined and randomly (and independently) selected the set of jobs each provider can perform. Hence, at 100% symmetry, our instances satisfy Assumption 1(ii) and are the same as those analyzed. In Figure B.5, we can see the average relative loss as a function of the number of providers, for different values of s (with $w = 15$, the region as Midtown Manhattan, and no geographical restrictions). The results suggest that for “high symmetry” (i.e., values of s above 60), the resulting average losses are not significantly different from those under full symmetry, recorded in Figures B.2 and B.3; but after instances become “sufficiently asymmetric” (the symmetry parameter drops below 50), the average losses decrease. Intuitively, lower values of s imply lower probabilities of two providers being able to perform the same job, which may explain why for low number of providers and low values of s , the loss is always zero: if no job can be performed by two providers, then the allocation problems can be separated into disjoint problems for each provider, which by Proposition 5, implies that the value loss will be zero. This shows that asymmetry can play a nuanced role, depending on other problem parameters: some asymmetry actually reduced the average losses here, but complete asymmetry also led to worst-case losses that asymptotically approached 100%, as seen in Instance A.4.

B.2. Synthetically Generated Instances

In order to analyze the dependency between the variation of values and the relative value loss we mention in Section 4, we generate synthetic instances where we can control this variation. In particular, for each instance we sample uniformly 30 starting points in the interval $(0, x] \subseteq \mathbb{R}$, for different x in the interval $(1, 3)$, and for each point we sample from a truncated normal distribution the length of the interval. We take this truncated normal distribution with mean 1 and a coefficient of variation cv varying from 0.001 to 0.5. Each interval represents a specific job (similarly to the trips in the Taxi and Limousine Commission (TLC) data). We consider the value of the each job to be exactly the length of the interval. As was the case with the trips, we will assume that two jobs cannot be allocated together if their intervals overlap. Therefore, we can measure the average loss for different values of the coefficient of variation of the intervals lengths. We take the average loss over 100 samples for each coefficient of variation. The results for a representative subset of values of x can be seen in Figure B.6. We observe that the main characteristics of the relative loss as a function of the ratio of providers to jobs is maintained for different values of x , with the difference that for lower x , we observe higher maximum values of relative loss. As occurred with the instances generated with the TLC data when lowering w , this may be because of the fact that we

Figure B.2. (Color online) Relative Value Loss for Different Regions

Note. Average and maximum $L_{\gamma}(\mathcal{F}, \mathcal{F}_G)$ as a function of the number of providers, for instances with 30 jobs constructed from the data, using $w = 20$ and three different regions of New York City.

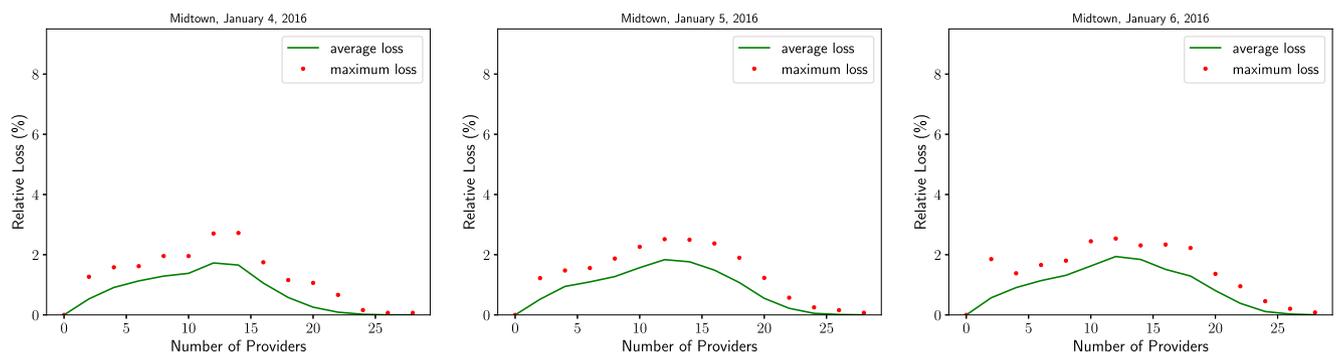
take instances of 30 jobs for all the values of x , which implies that the probability that 2 jobs are incompatible is lower for larger x .

We can see in Figure B.6, as we mentioned in Section 5, that the maximum loss is decreasing in the coefficient of variation of the values, consistent with the remarks of Section 4 on the variation of values as a driver of loss. Nevertheless, in the instances we generated for Figure B.6, the variation of the values is intrinsically connected to the variation in the lengths of the intervals we took to generate the feasibility restrictions. Hence, in order to isolate the effect of the variation of values, we took the same instances, but where we fixed the length of each interval (representing a job) to be exactly one, for feasibility purposes. Therefore, although the intervals that define the feasibility constraints all have length 1, the values remain as before, taken from truncated normal distributions with different coefficients of variation. We then plotted the average losses as before

in Figure B.7. By comparing Figure B.6 and B.7, we can see that the effect of the variation in values on the loss remains, although we do observe slightly higher losses, in particular for the cases with large coefficients of variation.

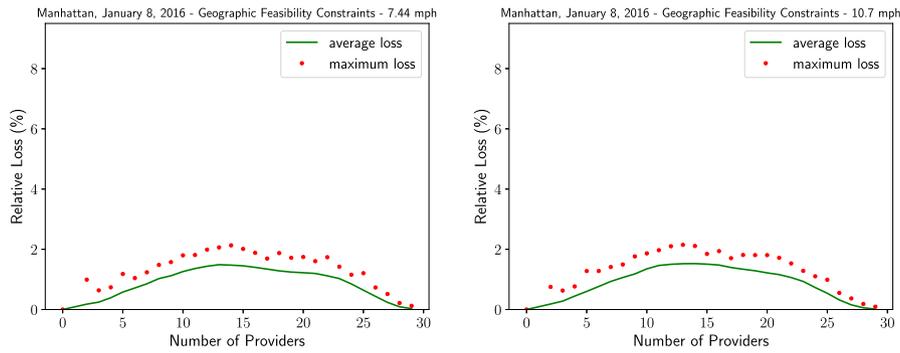
As a second robustness test on this effect we computed the loss for each instance when we completely remove the variation in values. For this, we simply take all instances we generated (both from the data and the synthetically generated), and we fix all values to be one. The resulting losses, under max-min fair guarantees, are always zero, for all instances. This once again affirms the importance of the variation in values as a main driver of loss in these instances, so much so that when we remove it, the losses disappear.

This numerical analysis shows that the average loss may be small in particular instances that are included in our general theoretical analysis. Moreover, it demonstrates the effect of many of the main drivers of loss we analyzed in Section 4.

Figure B.3. (Color online) Relative Value Loss for Different Dates

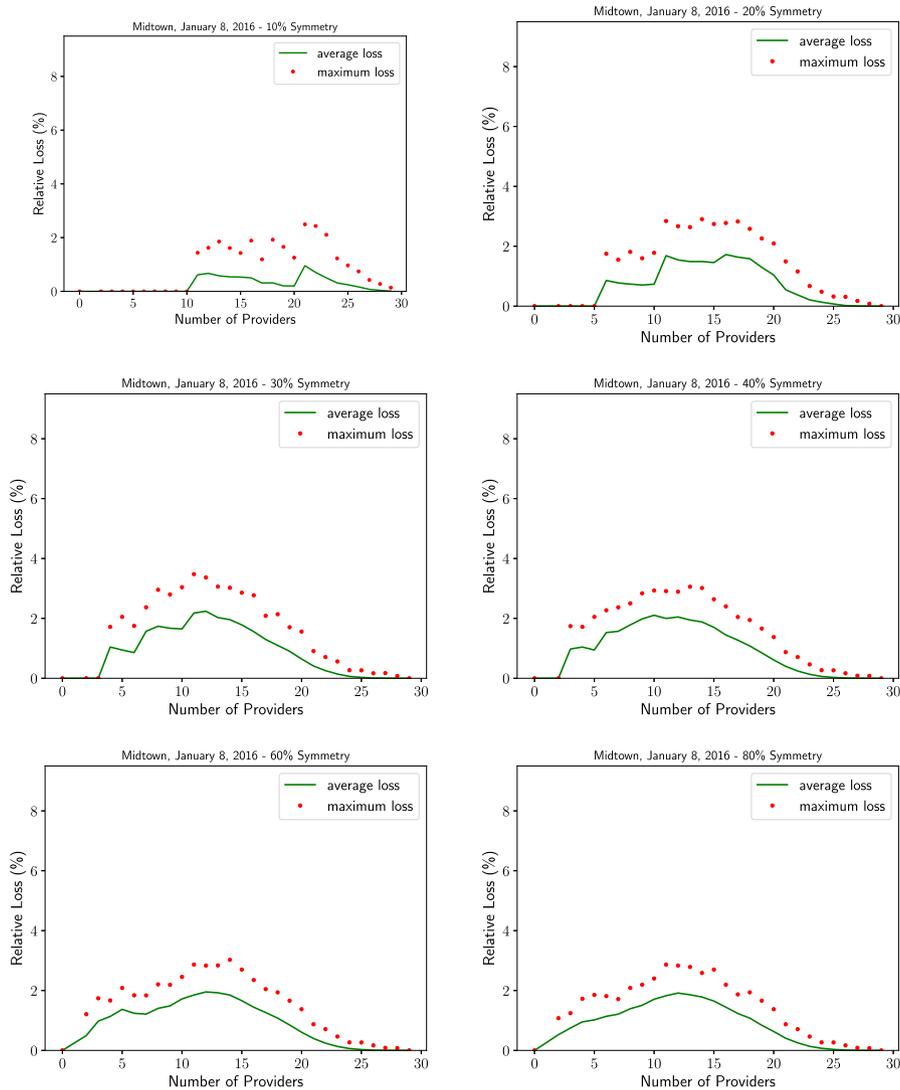
Note. Average and maximum $L_{\gamma}(\mathcal{F}, \mathcal{F}_G)$ as a function of the number of providers, for instances with 30 jobs constructed from the data, using $w = 15$ and three different dates of the first week of 2016.

Figure B.4. (Color online) Relative Value Loss with Geographic Constraints for Different Average Speed, v



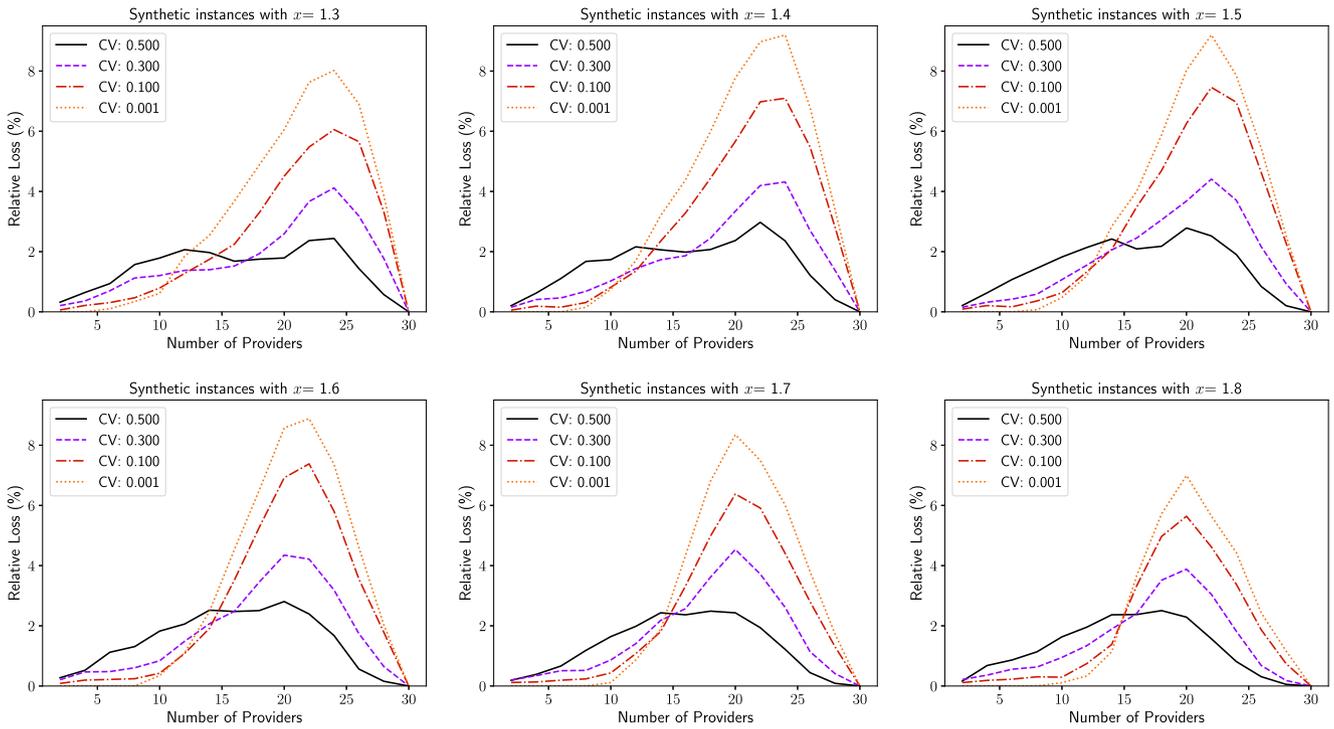
Note. Average and maximum $L_{\gamma}(\mathcal{F}, \mathcal{F}_G)$ as a function of the number of providers, for instances with 30 jobs constructed from the data with geographic constraints on the feasible allocations, using $w = 15$ and the two values of v : $v = 7.44$ (left panel) and $v = 10.44$ (right panel).

Figure B.5. (Color online) Relative Value Loss for Different Values of the Symmetry Parameter, s



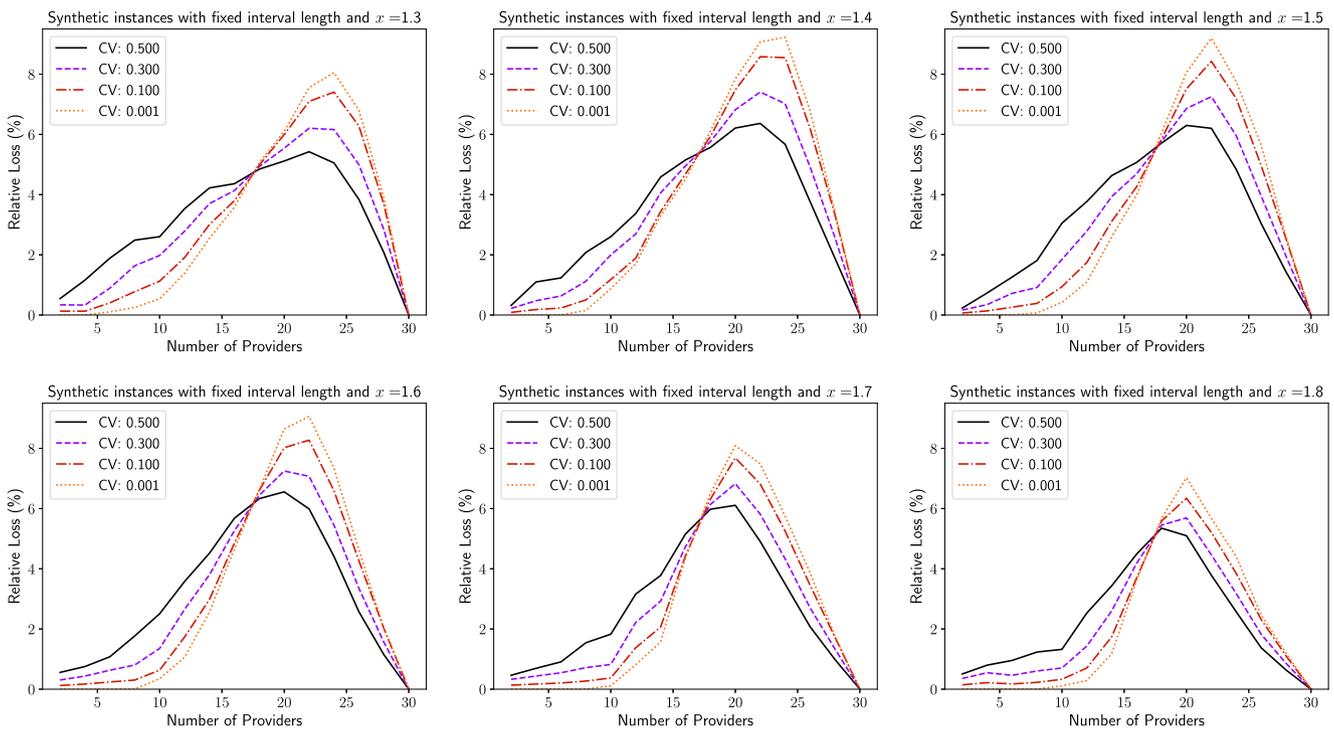
Note. Average and maximum $L_{\gamma}(\mathcal{F}, \mathcal{F}_G)$ as a function of the number of providers, for instances with 30 jobs constructed from the data with varying levels of symmetry s , using $w = 15$ and the six values of s : $s = 10\%$ (top left panel), $s = 20\%$ (top right panel), $s = 30\%$ (middle left panel), $s = 40\%$ (middle right panel), $s = 60\%$ (bottom left panel), and $s = 80\%$ (bottom right panel).

Figure B.6. (Color online) Synthetic Instances for x Varying from 1.3 to 1.8



Note. CV, coefficient of variation.

Figure B.7. (Color online) Synthetic Instances with Fixed Interval Length for x Varying from 1.3 to 1.8



Note. CV, coefficient of variation.

Endnotes

- ¹ In particular, the system could choose any allocation that maximizes the total value, and then redistribute this value through monetary transfers to ensure the provider guarantees are satisfied.
- ² For example, monetary transfers might lead to inequitable payment for identical jobs, which is linked to perceptions of unfairness (see, e.g., Greenberg 1982 and Brockner and Wiesenfeld 1996). In addition, it is unclear that monetary transfers can entirely mitigate nonmonetary aspects of the allocation and provider welfare.
- ³ Nevertheless, we show numerically that the effect of asymmetric sets of feasible allocations is nuanced: some asymmetry may actually reduce average losses, but complete asymmetry can lead to unbounded worst-case losses.
- ⁴ The same framework and analysis go through when considering a stochastic heterogeneity parameter γ_i with support $[\gamma_{\min}, \gamma_{\max}]$ and a known distribution, if we measure relative losses in the expected total value.
- ⁵ In particular, $u(\mathbf{B}) \geq u(\mathbf{A})$ for any $\mathbf{A}, \mathbf{B} \in \mathcal{F}$ with $v(B_i) \geq v(A_i), \forall i \in N$.
- ⁶ Note that by symmetry, because $(\emptyset, D) \in \mathcal{F}$, we must also have $(D, \emptyset) \in \mathcal{F}$, so this allocation is also feasible.
- ⁷ To provide one concrete example, Uber provides guarantees to new drivers (see, e.g., Rideshare Central 2018) who may perhaps be less productive than more experienced drivers (see, e.g., The Rideshare Guy 2016).
- ⁸ The max-min fair allocations are $\mathcal{F}_G^{\text{MM}} = \{(\{d_1\}, \{d_2\}), (\{d_2\}, \{d_1\})\}$. All such allocations do not assign d_3 because doing so would mean that one provider would obtain $\min_i p(A_i) = p(\{d_3\}) < p(\{d_1\})$, where the last inequality follows from the strict monotonicity of p and the fact that $v(d_3) = 1 - \kappa < 1 = v(d_1)$.
- ⁹ As in Instance 2, the max-min fair allocations are all combinations of $\{d_1, d_2, d_3, d_4\}$ into two sets of two jobs each, and all such allocations do not utilize d_5 .

References

Asadpour A, Saberi A (2010) An approximation algorithm for max-min fair allocation of indivisible goods. *SIAM J. Comput.* 39(7):2970–2989.

Atkinson A (1970) On the measurement of inequality. *J. Econom. Theory* 2(3):244–263.

Awerbuch B, Azar Y, Richter Y, Tsur D (2006) Tradeoffs in worst-case equilibria. *Theoretical Comput. Sci.* 361(2):200–209.

Bacharach SB, Bamberger P, Conley S (1991) Work-home conflict among nurses and engineers: Mediating the impact of role stress on burnout and satisfaction at work. *J. Organ. Behav.* 12(1):39–53.

Berger PD (1972) On setting optimal sales commissions. *J. Oper. Res. Soc.* 23(2):213–215.

Bertsekas D, Nedić A, Ozdaglar A (2003) *Convex Analysis and Optimization*, Athena Scientific Optimization and Computation Series (Athena Scientific, Nashua, NH).

Bertsimas D, Farias VF, Trichakis N (2011) The price of fairness. *Oper. Res.* 59(1):17–31.

Bertsimas D, Farias VF, Trichakis N (2012) On the efficiency-fairness trade-off. *Management Sci.* 58(12):2234–2250.

Bertsimas D, Farias VF, Trichakis N (2013) Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Oper. Res.* 61(1):73–87.

Besbes O, Iancu DA, Trichakis N (2018) Dynamic pricing under debt: Spiraling distortions and efficiency losses. *Management Sci.* 64(10):4572–4589.

Brockner J, Wiesenfeld BM (1996) An integrative framework for explaining reactions to decisions: Interactive effects of outcomes and procedures. *Psych. Bull.* 120(2):189–208.

Cachon GP, Lariviere MA (2005) Supply chain coordination with revenue-sharing contracts: Strengths and limitations. *Management Sci.* 51(1):30–44.

Chawla S, Roughgarden T (2008) Bertrand competition in networks. Monien B, Schroeder UP, eds. *Algorithmic Game Theory*, Lecture Notes in Computer Science, vol. 4997 (Springer, Berlin), 70–82.

CNBC (2017) Only 4% of Uber drivers remain on the platform a year later, says report. Accessed July 7, 2018, <https://www.cnn.com/2017/04/20/only-4-percent-of-uber-drivers-remain-after-a-year-says-report.html>.

Cohen-Charash Y, Spector PE (2001) The role of justice in organizations: A meta-analysis. *Organ. Behav. Human Decision Processes* 86(2):278–321.

Correa JR, Schulz AS, Stier-Moses NE (2004) Selfish routing in capacitated networks. *Math. Oper. Res.* 29(4):961–976.

DoorDash (2019) Challenges. Accessed November 14, 2019, https://help.doordash.com/dashers/s/challenges?language=en_US.

Eisenhardt KM (1988) Agency- and institutional-theory explanations: The case of retail sales compensation. *Acad. Management J.* 31(3):488–511.

Farley JU (1964) An optimal plan for salesmen's compensation. *J. Marketing Res.* 1(2):39–43.

Financial Times (2014) Uber offers US drivers perks to boost retention. Accessed July 7, 2018, <https://www.ft.com/content/446dd780-6f91-11e4-8d86-00144feabdcd>.

Gibbons R (1998) Incentives in organizations. *J. Econom. Perspect.* 12(4):115–132.

Golovin D (2005) Max-min fair allocation of indivisible goods. Technical report, Carnegie Mellon University, Pittsburgh.

Greenberg J (1982) Approaching equity and avoiding inequity in groups and organizations. Greenberg J, Cohen RL, eds. *Equity and Justice in Social Behavior* (Academic Press, New York), 389–435.

GrubHub (2019) Guaranteed hourly minimums. Accessed November 14, 2019, <https://driver-support.grubhub.com/hc/en-us/sections/360004883032-Guaranteed-Hourly-Minimums>.

Iancu DA, Trichakis N (2014) Fairness and efficiency in multiportfolio optimization. *Oper. Res.* 62(6):1285–1301.

Jensen MC, Meckling WH (1976) Theory of the firm: Managerial behavior, agency costs and ownership structure. *J. Financial Econom.* 3(4):305–360.

Johari R, Tsitsiklis JN (2004) Efficiency loss in a network resource allocation game. *Math. Oper. Res.* 29(3):407–435.

Kalai E, Smorodinsky M (1975) Other solutions to nash's bargaining problem. *Econometrica* 43(3):513–518.

Koutsoupias E, Papadimitriou C (2009) Worst-case equilibria. *Comput. Sci. Rev.* 3(2):65–69.

Lazear EP (2000) Performance pay and productivity. *Amer. Econom. Rev.* 90(5):1346–1361.

Le VB, Peng SL (2015) Characterizing and recognizing probe block graphs. *Theoretical Comput. Sci.* 568:97–102.

Lipton RJ, Markakis E, Mossel E, Saberi A (2004) On approximately fair allocations of indivisible goods. *Proc. 5th ACM Conf. Electronic Commerce—EC '04* (Aamas), 125.

Lyft Inc. (2018) Earnings guarantee promotions. Accessed July 7, 2018, <https://help.lyft.com/hc/en-ca/articles/115012927247-Weekly-Guarantees>.

Mas-Colell A, Whinston M, Green J (1995) *Microeconomic Theory* (Oxford University Press, Oxford, United Kingdom).

McCoy JH, Lee HL (2014) Using fairness models to improve equity in health delivery fleet management. *Production Oper. Management* 23(6):965–977.

- NYC Department of Transportation (2018) New York City Mobility Report June 2018. Accessed February 10, 2019, <http://www.nyc.gov/html/dot/downloads/pdf/mobility-report-2018-print.pdf>.
- NYC Taxi and Limousine Commission (2016) TLC trip record data. Accessed November 10, 2017, <https://www1.nyc.gov/site/tlc/about/tlc-trip-record-data.page>.
- Papadimitriou C (2001) Algorithms, games, and the internet. *Proc. 33rd Annual ACM Sympos. Theory Comput., STOC' 01* (ACM, New York), 749–753.
- Patagonia (2018) How we work with factories, farms and mills. Accessed July 27, 2018, <https://www.patagonia.com/working-with-factories-farms-mills/>.
- Perakis G, Roels G (2007) The price of anarchy in supply chains: Quantifying the efficiency of price-only contracts. *Management Sci.* 53(8):1249–1268.
- Qi J (2017) Mitigating delays and unfairness in appointment systems. *Management Sci.* 63(2):566–583.
- Rawls J (1971) *A Theory of Justice* (Harvard University Press, Cambridge, MA).
- Rideshare Central (2018) Uber sign up bonus. Accessed February 25, 2019, <https://web.archive.org/web/20190501002455/https://ridesharecentral.com/uber-sign-up-bonus>
- Roughgarden T (2003) The price of anarchy is independent of the network topology. *J. Comput. System Sci.* 67(2):341–364.
- Roughgarden T, Tardos E (2002) How bad is selfish routing? *J. ACM* 49(2):236–259.
- Simon HA (1956) Rational choice and the structure of the environment. *Psych. Rev.* 63(2):129–138.
- Sparks K, Cooper C, Fried Y, Shirom A (2011) The effects of hours of work on health: A meta-analytic review. *J. Occupational Organ. Psych.* 70(4):391–408.
- Starbucks Corporation (2018) Starbucks ethical sourcing of sustainable products. Accessed July 27, 2018, <https://www.starbucks.com/responsibility/sourcing>.
- The Rideshare Guy (2016) Why do so many people quit driving for Uber? Accessed February 25, 2019, <https://therideshareguy.com/why-do-so-many-people-quit-driving-for-uber/>.
- Tremblay M, Sire B, Balkin DB (2000) The role of organizational justice in pay and employee benefit satisfaction, and its effects on work attitudes. *Group Organ. Management* 25(3):269–290.
- Uber Technologies Inc. (2018a) Driver-partner incentives. Accessed February 10, 2020, <https://www.uber.com/drive/atlanta/resources/driver-partner-incentive-guarantee-faq-questions/>.
- Uber Technologies Inc. (2018b) Hourly guarantees. Accessed February 10, 2020, <https://www.uber.com/en-MY/drive/resources/earn-with-hourly-guarantees/>.
- Yang ZB, Aydın G, Babich V, Beil DR (2012) Using a dual-sourcing option in the presence of asymmetric information about supplier reliability: Competition vs. diversification. *Manufacturing Service Oper. Management* 14(2):202–217.
- Yu H, Zeng AZ, Zhao L (2009) Single or dual sourcing: Decision-making in the presence of supply chain disruption risks. *Omega* 37(4):788–800.