Value Loss in Allocation Systems
with Provider Guarantees

Yonatan Gur

Stanford University

Dan Iancu

Stanford University

Xavier Warnes[*]

Stanford University

Abstract

Centralized planning systems routinely allocate tasks to workers or service providers in order to generate the maximum possible value. These allocations can also critically influence the service providers' well-being, and thus the planning systems are often concerned with ensuring that their allocations satisfy particular desirable attributes. In some cases, such guarantees may induce a loss in the total value created or in the system's share of that value. We provide a broad framework that allows quantifying the magnitude of such value losses due to provider guarantees. We consider a general class of guarantees that includes many considerations of practical interest arising in the design of sustainable two-sided markets, workforce welfare and compensation, 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; but when few providers are present, the main loss driver is the heterogeneity in the providers' ability to generate value. We study additional loss drivers such as the variation in the value of jobs and the supply-demand balance, and find that these can influence the loss in a non-intuitive fashion, with less variability in value and a more balanced supply-demand ratio leading to larger losses. Lastly, we demonstrate numerically using both real-world and synthetic data that the relative loss can be very small in several cases of practical interest.

Keywords: Allocation systems, worker guarantees, fairness, efficiency loss, worst-case analysis.

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

---

[*]Correspondence: `ygur@stanford.edu`, `daniancu@stanford.edu`, `xwarnes@stanford.edu`.

shared between the system and the providers. And 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., due to 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 Uber, Lyft, Grubhub or Upwork. These platforms match customer service requests for rides, food, or labor with dedicated service providers (drivers or freelance workers), 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 non-monetary 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., Uber Technologies Inc. 2018b and Lyft Inc. 2018) to designing loyalty and bonus payments tied to completing multiple service requests (see, e.g., Financial Times 2014, and Uber Technologies Inc. 2018a). How much are platforms sacrificing in terms of revenues, profits, or rider experience when their allocations are designed to carry such guarantees for the service providers?

To demonstrate the tradeoff 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 potentially also customer satisfaction. But such allocations also have a direct impact on the employees, in both monetary terms, e.g., due to commissions and bonuses tied to completed sales (Berger 1972), as well as non-monetary ones, e.g., due to 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 due to 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 magnitude of value that may be lost due to imposing provider guarantees in various settings. We consider a centralized planning system that allocates a discrete set of jobs/resources to a set of heterogeneous providers who convert these into value. 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 very broad monotonicity condition. This allows us to capture a variety of practical considerations pertaining to sustainable design of two-sided markets, workforce welfare and compensation mechanisms, as well as the 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 out of the maximal value (captured by the unrestricted value-maximizing allocation) that

---

[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 non-monetary aspects of the allocation and provider welfare.

would be lost when imposing the constraints associated with satisfying these guarantees. We derive tight bounds on the relative value loss for *any* restrictions in a general class of provider guarantees. We establish these tight bounds by solving a fractional linear relaxation of the problem of maximizing the relative value loss, and by producing instances that match the maximal value under this relaxation. These bounds determine that the value loss cannot be arbitrarily large for any guarantees that we consider. Moreover, the relative value loss never exceeds $\frac{1}{2}$ when providers are homogeneous in their ability to generate value from jobs (which is the case, for example, in standardized work processes). This analysis enables us to characterize prominent drivers of the relative value loss. We show that when the number of providers is large, the largest value losses occur due to *fairness considerations* – and more precisely, Max-Min fairness. In contrast, when only few providers exist, loss is driven primarily by the heterogeneity in the providers' ability to generate value from the jobs that are allocated to them.

We study the potential impact of several candidate drivers of value loss. First, we observe that the structure of the set of feasible allocations may have a critical impact on the relative value loss, and we characterize several structures where this loss can be guaranteed to be zero. In addition, we also 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 in limiting the maximal value loss: when providers are heterogenous with respect to their ability to create value from jobs, the relative value loss may asymptotically approach 100% as the number of providers grows large. We further demonstrate that the variation in the intrinsic values associated with the jobs impacts the value loss: in general, higher variation in these values induces larger worst-case losses. Finally, we show that an imbalance between supply (providers) and demand (jobs) may lead to reduced value losses.

Using both real-world data of taxi trips in New York City (where providers could be viewed as approximately homogeneous with respect to their ability to generate value from jobs) and synthetically generated data, we study numerically the effects of the aforementioned drivers of loss on the relative value loss associated with implementing provider guarantees that correspond to fairness considerations. This confirms the robustness of our earlier results concerning relationships between the relative value loss and the variation in the intrinsic values of jobs, as well as the imbalance between providers and jobs. We observe, in particular, that in the instances generated

from the real-world data, the relative value loss that corresponds with fairness considerations remains below 4%. Together with the rest of our numerical findings, this suggests that in particular cases of practical interest, the relative value losses associated with implementing provider guarantees may be low relative to the worst-case value loss bounds characterized here.

1.2 Related Literature

Relative Efficiency Loss. Our work is related to a stream of literature that studies efficiency loss when some feasible outcomes are restricted. 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 provide bounds on this measure that depend on the number of agents and on the fairness criterion. In contrast, our study focuses on discrete allocation problems where constraints are imposed directly on the feasible allocations rather than on the possible utility outcomes. The class of restrictions we consider is thus broader and includes the fairness criteria in Bertsimas et al. (2011, 2012) as special cases. Additionally, while the worst-case efficiency loss in Bertsimas et al. (2011, 2012) can be arbitrarily large (i.e., approaching 100% asymptotically), the loss in our setting is always bounded for any fixed provider heterogeneity level. We elaborate more on the root causes for this discrepancy in §3.

More broadly our paper relates 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 particular settings (see, e.g., Roughgarden 2003, Johari and Tsitsiklis 2004, Correa et al. 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). Our study considers efficiency losses generated when a central organizer restricts the outcomes, rather than losses resulting from the actions of selfish agents.

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; in contrast, our paper is focused on quantifying the inefficiency associated

5

with such allocations, and studying the key drivers of this inefficiency.

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), among other examples. Several papers in this literature also seek to quantify the associated efficiency losses; see, e.g., Besbes et al. 2017 for more discussion and additional references. Our work is not concerned with specific agency considerations, but instead focuses on quantifying 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 §2.1. A discussion of the modeling assumptions is deferred to §2.2.

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 $v : D \to \mathbb{R}$, 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 $\boldsymbol{A} = (A_1, \dots, A_n) \in \mathcal{F}$ 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 pre-determined and fixed, and belongs to an interval $[\gamma_{\min}, \gamma_{\max}]$, where $0 < \gamma_{\min} \leq \gamma_{\max}$[3]. We let $\boldsymbol{\gamma} \in [\gamma_{\min}, \gamma_{\max}]^n$ denote the heterogeneity profile of providers, and we denote the degree of heterogeneity by

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

---

[3]The same framework and analysis goes 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.

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

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

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* $\mathcal{F}_{\mathrm{G}} \subseteq \mathcal{F}$. This could reduce the total value generated, and we define the *value loss under provider guarantees* $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}})$ as the relative loss in total value when the system considers only allocations from $\mathcal{F}_{\mathrm{G}}$:

$$\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) = \frac{\displaystyle\max_{\boldsymbol{A} \in \mathcal{F}} \sum_{i=1}^{n} \gamma_i v(A_i) - \max_{\boldsymbol{B} \in \mathcal{F}_{\mathrm{G}}} \sum_{i=1}^{n} \gamma_i v(B_i)}{\displaystyle\max_{\boldsymbol{A} \in \mathcal{F}} \sum_{i=1}^{n} \gamma_i v(A_i)}. \tag{1}$$

Throughout the paper, we restrict attention to cases where $\mathcal{F}_{\mathrm{G}}$ is non-empty. 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.

2.1 Examples

We next illustrate the interpretation of various model components using a series of practical examples.

Service platforms. Service platforms such as Uber and Lyft can be thought of as systems that allocate demand for services (i.e., trips) to providers (i.e., drivers). The trips are indivisible, and each trip has a certain intrinsic value $v(d)$, with several possible interpretations. For instance, $v(d)$ may correspond to the revenue from completing trip 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 trips, a measure of economic efficiency in the absence of costs.

(ii) Revenue sharing. It is customary for service platforms such as Uber, Lyft or Grubhub 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 length of a trip are typically key drivers for both the revenues as well as the costs from completing the trip. 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, due to 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 trips, 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 trip d, then $\gamma_i v(d)$ could also capture the quality of service experienced by the rider(s) when driver i completes the trip. Under this interpretation, γ_i may capture idiosyncratic differences due to car cleanliness, driver friendliness, or any other feature that influences the rider's experience per unit of trip length. The total generated value would then correspond to the total quality of service experienced by riders from the allocations. The set of feasible allocations $\mathcal{F}$ could capture constraints on allocating trips 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, for example, Uber Technologies Inc. 2018b). In this setting, the value loss would thus be driven by two key factors: drivers' heterogeneity in their ability to generate value (monetary or non-monetary) from trips allocated to them, as well as potential demand loss, e.g., from the inability to allocate all requests for trips.

Workforce Management. A closely related example arises 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 ability/performance), or it could capture the fraction of the revenue accruing to the firm (with $1 - \gamma_i$ denoting associate i's commission). Similar to 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).

Sourcing from a Heterogeneous Supply Base. A different example arises when the system represents a firm that decides how to allocate pre-scheduled 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 or quality), gross margins for the firm (when different prices are paid to different suppliers), or gross margins for the entire supply chain (for example, when suppliers have different variable operating costs). 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).

2.2 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, i.e.,

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

ii) (Symmetry) If $\boldsymbol{A}$ is a feasible allocation, then any permutation of $\boldsymbol{A}$ 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 $i \in N$.

9

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 $(B_1, \ldots, B_{i-1}, A_i, B_{i+1}, \ldots, B_n) \in \mathcal{F}$.

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 (i) service platforms such as Uber, Lyft, Grubhub or Upwork, where unique jobs must be allocated to service providers operating independently; (ii) sales settings where a single lead cannot be divided across multiple associates; and (iii) sourcing settings where a single unit cannot be divided among multiple 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 Uber/Lyft drivers may all be able to complete a particular ride, but the value that is generated (e.g., the profit or the quality of service experienced by the rider) 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 concerns the set of allocations with guarantees $\mathcal{F}_G$.

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

Assumption 2 provides a connection between the value that is captured by completing a job, and the allowable provider guarantees that could be considered. Broadly speaking, valid restrictions correspond to provider guarantees that compensate for additional value: if a certain allocation satisfies the required provider guarantees, then any feasible allocation in which each provider generates more value should also satisfy these guarantees. As we demonstrate in the following subsection, such an alignment between captured value and provider guarantees is natural in many practical settings.

2.3 Discussion and Classes of Provider Guarantees

2.3.1 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 providers when compensation to providers is increasing in the value they capture 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 $p : \mathcal{P}(D) \to \mathbb{R}$ denotes a payment function. For a given real number τ, we define the set

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

that 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:

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

The latter requirement is a natural property for payment or compensation functions: it asks that a job carrying more intrinsic value should also command a (weakly) higher compensation when completed. An important family of payment functions that satisfy this property is the one of proportional compensation functions of the form $p(S) = \theta v(S)$ for some $\theta \in [0, 1]$. Proportional payment functions are widely used in practical revenue-sharing systems (including Lyft, Uber, and Upwork, among many others) where service providers retain a constant fraction of the payment

11

made by the consumer once the job is completed. Proportional payment functions 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, i.e. $p(S) = \theta v(S)$ for $\theta = 1$. The figure shows both *uniform income guarantees* (with a unique τ for all providers), as well as non-uniform 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 $\mathcal{F}_\text{G}$ 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 §2.3.3.

To facilitate our subsequent analysis, we introduce a characterization of the set $\mathcal{F}_\text{G}$ that is equivalent to verifying Assumption 2.

Proposition 1 *A subset $\mathcal{F}_\text{G} \subseteq \mathcal{F}$ satisfies Assumption 2 if and only if $\mathcal{F}_\text{G}$ satisfies:*

$$\mathcal{F}_\text{G} = \arg\max_{\boldsymbol{A} \in \mathcal{F}} g(\boldsymbol{A}),$$

for some $g : \mathcal{F} \to \mathbb{R}$ that satisfies $g(\boldsymbol{B}) \geq g(\boldsymbol{A})$ for any $\boldsymbol{A}, \boldsymbol{B} \in \mathcal{F}$ 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 valid restrictions correspond to the allocations that would maximize some function g that preserves the same ordering as the value function v. One could now observe that condition (3) guarantees that Assumption 2 is satisfied, since the function $g(\boldsymbol{A}) \overset{\text{def}}{=} \mathbb{1}\{p(A_i) \geq \tau_i, \forall i \in N\}$ satisfies the condition in Proposition 1.

2.3.2 Unions and Intersections.

Consider any collection of sets $\{\mathcal{F}_G^k\}_{k \in K}$ where each set $\mathcal{F}_G^k$ satisfies Assumption 2. Then, $\cap_{k \in K} \mathcal{F}_G^k$ and $\cup_{k \in K} \mathcal{F}_G^k$ also satisfy Assumption 2. Considering intersections is useful for modeling restrictions that certain targeted providers find desirable or acceptable. Namely, suppose each provider $i \in N$ is endowed with a utility function u^i satisfying the co-monotonicity requirements in Proposition 1.[4] In

---

[4]In particular, $u(\boldsymbol{B}) \geq u(\boldsymbol{A})$ for any $\boldsymbol{A}, \boldsymbol{B} \in \mathcal{F}$ with $v(B_i) \geq v(A_i)$, $\forall i \in N$.
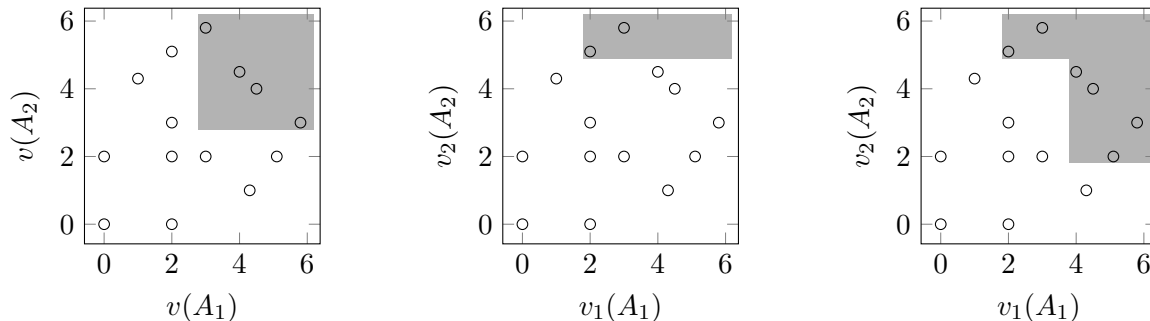
Figure 1: **Examples of provider guarantees.** Income guarantees $\mathcal{F}_G(p, \tau, N)$ for two providers ($n = 2$) compensated according to revenue-sharing agreements with $p(S) = \theta v(S)$ for $\theta = 1$. Each circle denotes a feasible allocation, with the two axes corresponding to the intrinsic 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)** A uniform income guarantee with $\tau = 3$; **(center)** A non-uniform income guarantee with $\tau_1 = 2$, $\tau_2 = 5$; **(right)** The union of two non-uniform income guarantees with $(\tau_1, \tau_2) = (2, 5)$ and $(\tau_1, \tau_2) = (4, 2)$, respectively.

addition, let

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

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

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, it can be shown that *any* provider guarantee satisfying Assumption 2 can actually be written as the union of income guarantees under monotonic payment functions, as formalized in our next result.

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 $\{(\boldsymbol{p}_k, \boldsymbol{\theta}_k)\}_{k \in K}$ for some index set K such that $\mathcal{F}_G = \cup_{k \in K} \mathcal{F}_G(\boldsymbol{p}_k, \boldsymbol{\theta}_k, N)$.*

This shows 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.

2.3.3 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:

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

Because this welfare function is increasing in each component A_i, the associated restriction $\mathcal{F}_\text{G} = \arg\max_{\boldsymbol{A}\in\mathcal{F}} W_\alpha(\boldsymbol{A})$ satisfies Assumption 2 (this is an immediate corollary of Proposition 1). A special case is *Max-Min fairness*, a concept inspired by the notion of Rawlsian justice (Rawls 1971). Max-Min fair allocations result from uniform income guarantees under monotonic payment functions, when the guarantee τ is the largest possible value for which the restriction set $\mathcal{F}_\text{G}$ is non-empty. More formally, we define the *Max-Min fair restriction under monotonic payments* as

$$
\mathcal{F}_\text{G}^\text{mM}(p, N) := \mathcal{F}_\text{G}(p, \tau_\text{max}, N), \qquad \text{where} \qquad \tau_\text{max} := \max\{\tau \mid \mathcal{F}_\text{G}(p, \tau, N) \neq \emptyset\}. \tag{4}
$$

3 Bounding the Value Loss under Provider Guarantees

Our first result provides an upper bound on the relative value loss $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_\text{G})$ 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:*

$$
\sup_{\mathcal{F}, \mathcal{F}_\text{G}, \boldsymbol{\gamma}} \mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_\text{G}) \leq \max\left\{ \delta, \frac{n-1}{n + (1-\delta)(n-1)} \right\}, \tag{5}
$$

where the supremum is taken with respect to all $\boldsymbol{\gamma} \in [\gamma_\text{min}, \gamma_\text{max}]^n$ and all sets $\mathcal{F}$, $\mathcal{F}_\text{G}$ 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.
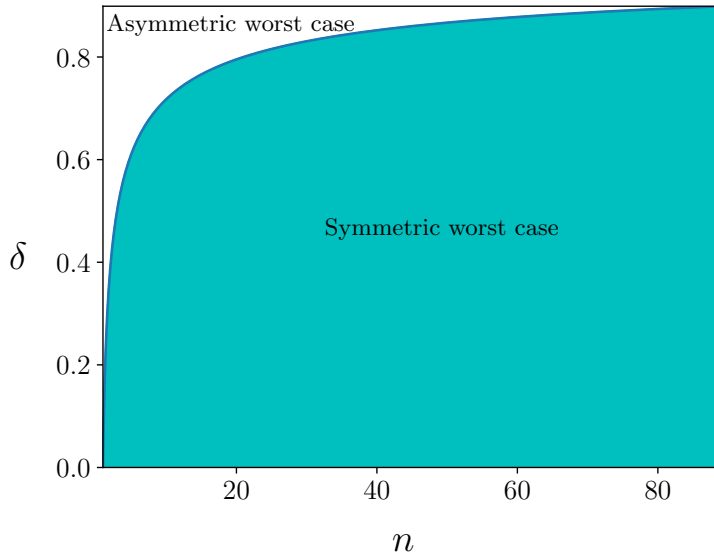

Figure 2: **Different types of worst cases in different parametric regions.** Shaded area denotes values of (n, δ) such that $\delta \leq \frac{n-1}{n+(1-\delta)(n-1)}$.

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 exceeds $\frac{1}{2}$ when providers are perfectly homogeneous ($\delta = 0$). Importantly, and as emphasized in §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).

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 an LP relaxation of the problem of maximizing $\mathbf{L}_{\gamma}(\mathcal{F}, \mathcal{F}_{\mathrm{G}})$ over all sets $\mathcal{F}$ and $\mathcal{F}_{\mathrm{G}}$, for a fixed vector γ. For this purpose, we provide a family of
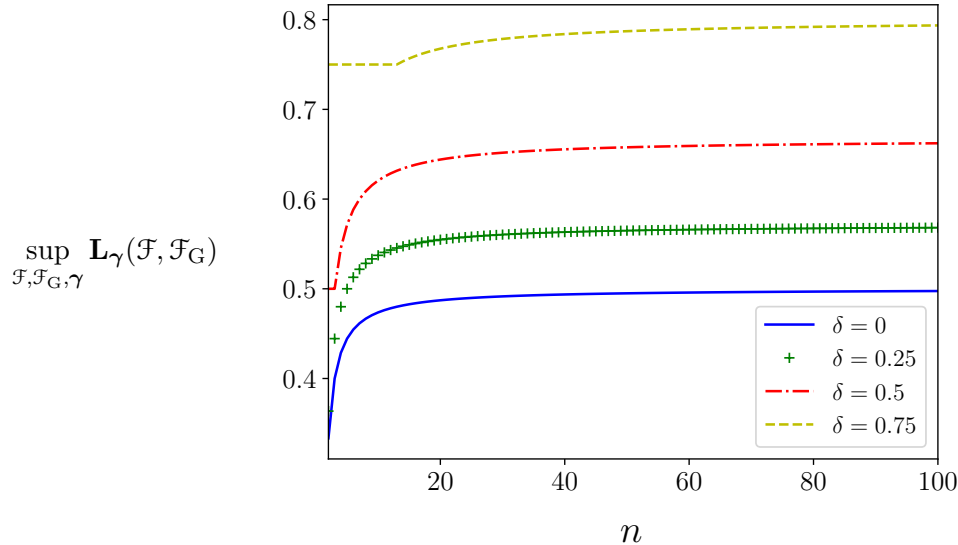
Figure 3: **Worst-case loss as a function of the number of providers and the heterogeneity level.** Plot of $\max\left\{\delta, \frac{n-1}{n+(1-\delta)(n-1)}\right\}$ as a function of n for different values of δ.

inequalities that connect the maximum value over the set of feasible allocations $\mathcal{F}$ and over the set of allocations with guarantees $\mathcal{F}_{\mathrm{G}}$. By exploiting the quasi-convexity of the optimal value of this relaxation as a function of $\boldsymbol{\gamma}$, we can then maximize the loss by only considering extreme heterogeneity profiles $\boldsymbol{\gamma} \in \{\gamma_{\min}, \gamma_{\max}\}^n$. Finally, we solve the LP relaxation for these values of $\boldsymbol{\gamma}$ 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 $\mathcal{F}^1, \mathcal{F}_{\mathrm{G}}{}^1, \boldsymbol{\gamma}^1$ and $\mathcal{F}^2, \mathcal{F}_{\mathrm{G}}{}^2, \boldsymbol{\gamma}^2$ satisfying Assumptions 1-2 so that*

$$\mathbb{L}_{\boldsymbol{\gamma}^1}(\mathcal{F}^1, \mathcal{F}_G^1) = \delta \tag{6}$$

$$\mathbb{L}_{\boldsymbol{\gamma}^2}(\mathcal{F}^2, \mathcal{F}_G^2) = \frac{n-1}{n+(1-\delta)(n-1)} - \varepsilon. \tag{7}$$

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.

16

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 $D = \{d_1\}$ with $v(d_1) = 1$, a set of feasible allocations $\mathcal{F} = \{(D, \emptyset), (\emptyset, D)\}$, and allocations with guarantees $\mathcal{F}_G = \{(\emptyset, D)\}$. Then, $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_G) = \frac{\gamma_{\max} - \gamma_{\min}}{\gamma_{\max}} = \delta$.*

In Instance 1, the system has a single job to assign to two providers. The first provider generates more value for the system ($\gamma_1 \geq \gamma_2$), so the value-maximizing allocation would be $(D, \emptyset)$, assigning the job to that provider. However, the only allocation with guarantees is $(\emptyset, D)$, requiring the job to be assigned to the second provider and thus generating a loss due to heterogeneity. It is important to note that these guarantees are asymmetric: a particular provider must be allocated all the jobs. In fact, if these guarantees had been "symmetrized," e.g., by considering the union of all permutations of which driver receives the full set of jobs, then we would have $\mathcal{F}_G = \mathcal{F}$, and 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 more jobs being assigned to less effective providers.

Another distinctive property of Instance 1 is that all the jobs are allocated, and hence the loss is only driven by the provider heterogeneity. This leads to worst-case losses when heterogeneity is high, as depicted in Figure 2. This inefficiency can take 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.[5]

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 non-overlapping. The allocations with guarantees are given by all Max-Min fair allocations $\mathcal{F}_G^{\mathrm{mM}}(p, N)$ under any strict monotonic payment function p, as described in (4). Then, the value-maximizing allocation is $\big(\{d_1, d_2\}, \{d_3\}\big)$, the value-maximizing*

---

[5]To provide one concrete example, Uber provides guarantees to new drivers (see, e.g., Rideshare Central 2018), that may perhaps be less productive than more experienced drivers (see, e.g., The Rideshare Guy 2016).

allocation with guarantees is $\big(\{d_1\}, \{d_2\}\big)$,[6] and the relative value loss is:

$$\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) = \frac{(2\gamma_{\max} + (1-\kappa)\gamma_{\min}) - (\gamma_{\max} + \gamma_{\min})}{2\gamma_{\max} + (1-\kappa)\gamma_{\min}}.$$

In particular, for any $\epsilon > 0$ there exists a $\kappa > 0$ such that $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) = \frac{1}{3-\delta} - \epsilon$, and thus

$$\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) \underset{\kappa \to 0}{\to} \frac{1}{3-\delta} = \frac{n-1}{n + (1-\delta)(n-1)}.$$
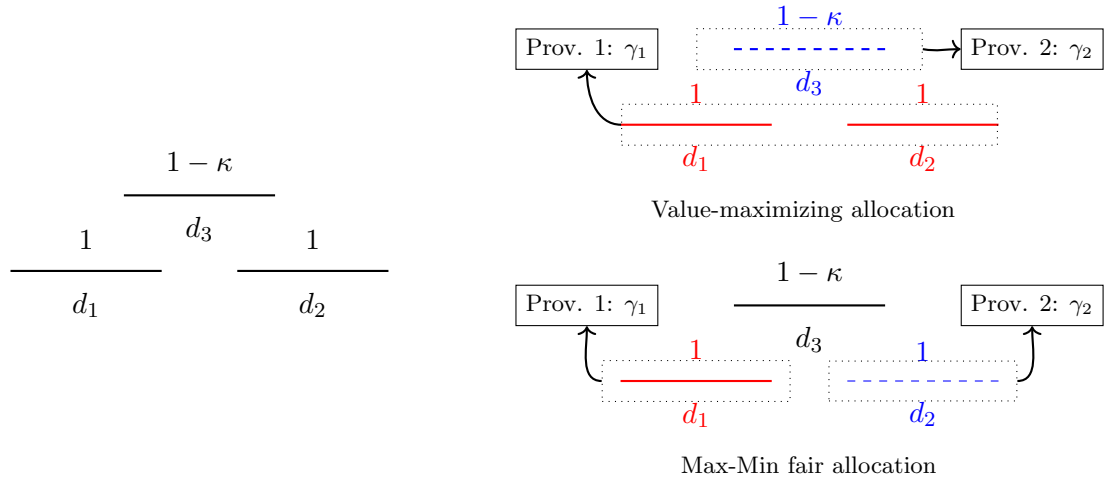

Figure 4: **Symmetric worst-case instance.** The left panel depicts Instance 2 that consists of $n = 2$ providers and three jobs $D = \{d_1, d_2, d_3\}$ represented by segments. Jobs can be assigned together if the corresponding segments are non-overlapping. The right panel depicts the value-maximizing allocation at the top, and the Max-Min fair allocation at the bottom.

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 since 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 $d_{1,2}$.

---

[6] The Max-Min fair allocations are $\mathcal{F}_{\mathrm{G}}^{\mathrm{mM}} = \big\{ \big(\{d_1\}, \{d_2\}\big), \big(\{d_2\}, \{d_1\}\big) \big\}$. All such allocations do not assign d_3, since 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 $v(d_3) = 1 - \kappa < 1 = v(d_1)$.

Instance 2 thus showcases two new critical 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 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 $\mathbf{L}_{\gamma}(\mathcal{F}, \mathcal{F}_G)$ would vanish. We revisit such exclusion constraints in §4, where we analyze their impact on the value loss in more detail. Second is the heterogeneity in the intrinsic value of jobs. Although the loss grows as jobs become more similar in value (i.e., as $\kappa \to 0$), some heterogeneity is in fact critical: if $\kappa = 0$, the loss $\mathbf{L}_{\gamma}(\mathcal{F}, \mathcal{F}_G)$ would again vanish. However, although this feature is important in the context of Instance 2, it is not necessary for achieving a worst-case loss in general, as our next instance demonstrates.

Instance 3 (Max-Min fairness with monotonic payments and equal-valued jobs) *Consider $n = 2$ providers with $\gamma_1 = \gamma_{\max} \geq \gamma_2 = \gamma_{\min}$, and a set of jobs $D = \{d_1, d_2, d_3, d_4, d_5\}$ 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 non-overlapping. 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 $\big(\{d_1, d_2, d_3, d_4\}, \{d_5\}\big)$, and a value-maximizing allocation with guarantees is $\big(\{d_1, d_2\}, \{d_3, d_4\}\big)$.[7] Therefore, $\mathbf{L}_{\gamma}(\mathcal{F}, \mathcal{F}_G) = \frac{4\gamma_{\max} + \gamma_{\min} - 2(\gamma_{\max} + \gamma_{\min})}{(4\gamma_{\max} + \gamma_{\min})}$.*
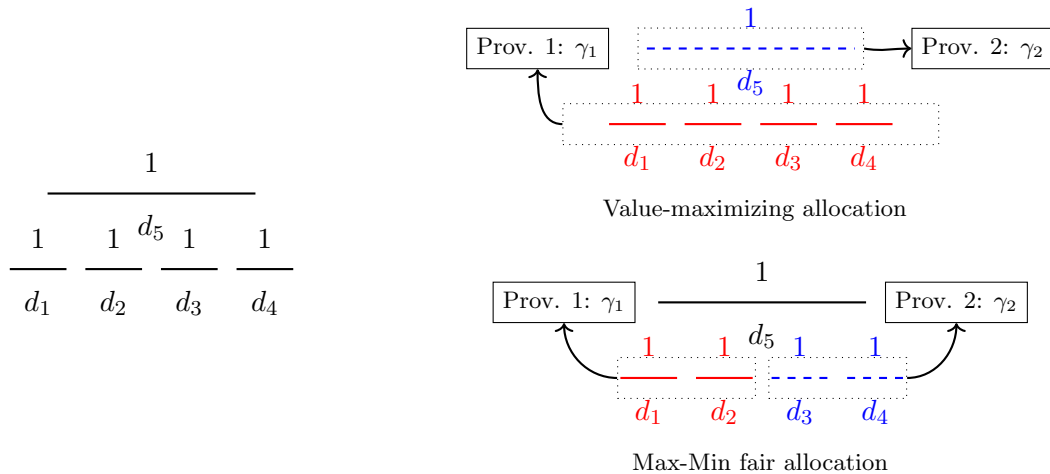

Figure 5: **Symmetric worst-case instance with equal-valued jobs.** The left panel depicts Instance 3: there are $n = 2$ providers and five jobs $D = \{d_1, d_2, d_3, d_4, d_5\}$ represented by segments. Jobs can be assigned together if the corresponding segments are non-overlapping. The right panel depicts the value-maximizing allocation at the top, and the value-maximizing Max-Min fair allocation at the bottom.

[7]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.

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 8 of Appendix A). This generalization yields a worst-case loss of:

$$
\begin{aligned}
\mathbf{L}_\gamma(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) &= \frac{(t+1)n\gamma_{\max} + t(n-1)\gamma_{\min} - (t+1)(\gamma_{\max} + (n-1)\gamma_{\min})}{(t+1)n\gamma_{\max} + t(n-1)\gamma_{\min}} \\
&= \frac{t(n-1) + \delta(n-1)}{t(n + (n-1)(1-\delta)) + n}.
\end{aligned}
$$

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

Instance 3 shares certain similarities with Instance 2. Both instances rely on the presence of exclusion constraints in the set of feasible allocations $\mathcal{F}$ that prevent certain jobs from being allocated together. Additionally, the worst-case guarantees in both instances correspond to symmetric (in particular, Max-Min fairness) guarantees obtained under any strictly monotonic payment function. This suggests that when the provider heterogeneity is not too large, symmetric guarantees (and perhaps particularly fairness) can be critical drivers of loss when there are sufficiently many providers, as depicted in Figure 2.

Instance 2 also exhibits two notable differences from Instance 3: 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 heterogeneity in job values and the supply/demand imbalance — and their relationship in more detail in §4. Lastly, it is worth emphasizing that although both Instance 2 and Instance 3 involve some heterogeneity in the drivers' value generation, this heterogeneity is not a critical to the identification of these instances as worst-case instances; in fact, both instances generate worst-case losses when providers are homogeneous ($\delta = 0$).

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 are very different in their ability to generate value (i.e., δ is large), this heterogeneity becomes the *dominant* loss driver. This is evidenced in Instance 1, and by the presence of the single term δ in the worst-case loss expression in (5). To isolate the effects, we thus focus our discussion henceforth on the case of homogeneous providers ($\delta = 0$). We explore several loss drivers: (i) the indivisibility of jobs; (ii) the structure of feasible

allocations $\mathcal{F}$; (iii) the variation in the intrinsic values of jobs; and (iv) the balance between supply (number of providers) and demand (number of jobs).

4.1 Indivisibility of Jobs

The indivisibility of jobs turns out to be critical: 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 $\mathcal{F}$:

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

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 $\boldsymbol{A}^j$. Hence, for $\boldsymbol{C} = (\{\theta_j\}_{j=1}^k, \{\boldsymbol{A}^j\}_{j=1}^k) \in \mathcal{F}^c$, 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

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

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

Proposition 3 *Given 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, $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}^c, \mathcal{F}_G^c) = 0$ for any set of uniform income guarantees under monotonic payments $\mathcal{F}_G^c$.*

The intuition behind Proposition 3 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.

4.2 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 become 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 these 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.

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 $\mathcal{F} = \{(A_1, \ldots, A_n) \mid A_i \subseteq D, |A_i| \leq k_i, \forall i, \text{ and } A_i \cap A_j = \emptyset, \forall i \neq j\}$.*

Proposition 4 *Instance 4 satisfies $\mathbf{L}_\gamma(\mathcal{F}, \mathcal{F}_G) = 0$ 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 a set of jobs $D = S \cup C$, and a set of feasible allocations of the form:*

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

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 would thus induce zero value 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 following result.

Proposition 5 *Consider Instance 5, then:*

(i). If $v(d) = v$, for all $d \in D$, and $|S| < 2n$, $\mathbf{L}_\gamma(\mathcal{F}, \mathcal{F}_G) = 0$ 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 5 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).

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 were relaxed (e.g., if providers had different abilities for completing jobs), then we can actually achieve a worst-case loss that asymptotically approaches 100% as the number of agents grows large (see Instance 9 of Appendix A). This instance is actually inspired by — and matches — the upper-bound on the price of fairness proved by Bertsimas et al. (2011), and shows that such asymmetries in the feasible sets of allocations can be a critical driver for the loss.

4.3   Variation in Intrinsic Job Values

Instance 3 showed that the difference in the intrinsic value of jobs can be a key driver of the loss. To further explore the impact of this feature, we now consider a slight modification of Instance 2 where we introduce a random variation in the value of one of the jobs, and we consider the expected loss as a function of this variation.

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:*

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

Note that $\mathbb{E}[\mathbf{L}_\gamma(\mathcal{F}, \mathcal{F}_G)]$ is decreasing in Δ; and since 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 a higher variation in values could actually *reduce* the impact of implementing provider guarantees, and result in a *lower* loss. The example can be

generalized to a case with n providers (see Example 2 in Appendix A), and we also confirm its robustness in a more realistic setting as part of our numerical exercise in §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 §5.

4.4   Supply-Demand Balance

Instance 3 suggested that the balance of supply (number of providers) and demand (number of jobs) can also critically drive the value loss. To build some further intuition for this dependency, 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 are present, the value loss vanishes.

Proposition 6 *Assume that $\gamma_{\min} = \gamma_{\max}$, and consider any symmetric 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 7 *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 above results suggest that the value loss is likely small when 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 §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.

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.

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 $n \in \{2, 3, \ldots, 30\}$. We assume that drivers are homogeneous, and thus $\gamma_{\min} = \gamma_{\max} = \gamma$, which we normalize without loss of generality to 1. 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 $(A_1, \ldots, A_n)$ 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 non-overlapping trips. The set of allocations with guarantees $\mathcal{F}_{\mathrm{G}}$ is obtained by considering Max-Min fairness considerations under revenue sharing, as discussed in §2.3.1. To obtain the relative value loss $\mathbf{L}_\gamma(\mathcal{F}, \mathcal{F}_{\mathrm{G}})$, we compute the value-maximizing allocation and the best restricted allocation by solving integer programming problems when allocations are picked from $\mathcal{F}$ and $\mathcal{F}_{\mathrm{G}}$, respectively. Further details on our setup are provided in Appendix B.

Real data. We generate problem instances using a publicly available dataset 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 location, 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 allows to more consistently suit part (*ii*) of Assumption 1 concerning the standardized nature of jobs, in that any provider in the area can perform any subset of (non-overlapping) trips. 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 9am and 5pm, 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 $w \in \{10, \ldots, 20\}$. 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.

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 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 1 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.

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 a smaller sets of jobs that can be allocated together.
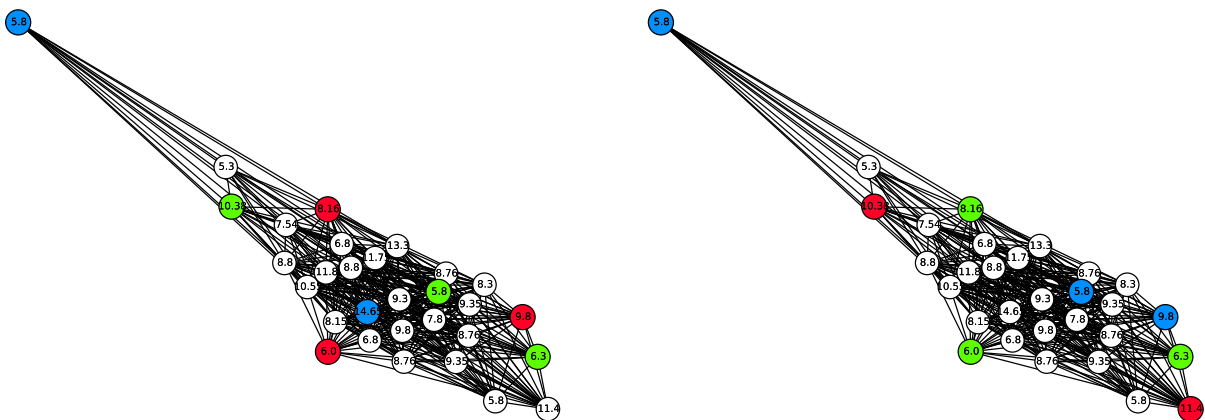

Figure 6: **Graph representation of the set of feasible allocations $\mathcal{F}$ of an instance.** Each node represents a trip, labels represent the fare, and two nodes are connected by an edge when they cannot be allocated together. **(left)** A value-maximizing allocation of the jobs to $n = 3$ providers, with the allocation to each provider given by a different color. The total value allocated is \$66.89, and the provider with the smallest allocation is receiving \$20.45. **(right)** A (value-maximizing) max-min fair allocation. The total value allocated is \$63.64, and the provider with the smallest allocation is receiving \$20.46.

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 the result in Proposition 5, and is reasonable to expect precisely because many of our data-driven instances match Instance 5, which
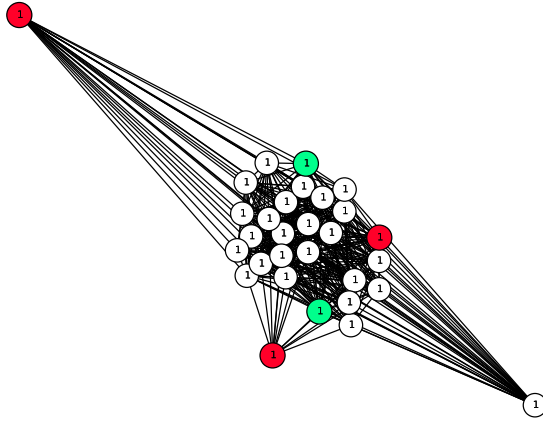
Figure 7: **Graph representation of $\mathcal{F}$ for an instance where trips have identical intrinsic values.** Each node represents a trip and two nodes are connected by an edge when they cannot be allocated together. A value-maximizing and Max-Min fair allocation with 2 providers is represented by the coloring. The maximum amount of trips that can be allocated to two providers is 5, by allocating the only three trips that can be completed together to one provider, and two other trips to the other provider.

is the premise for Proposition 5 (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 figure 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 with the results in Proposition 6, 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 zero when the number of providers
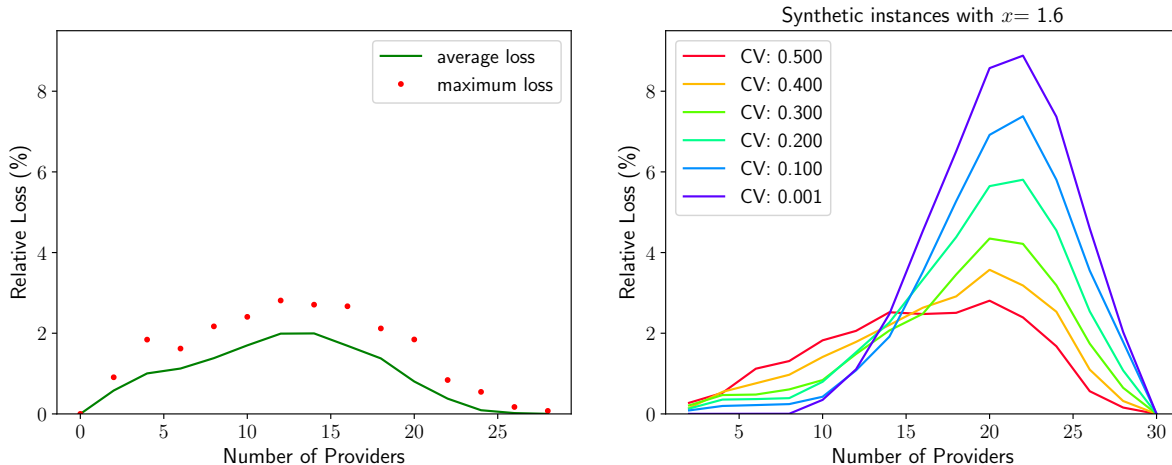
27

Figure 8: **Relative value loss in particular instances. (left)** Average and maximum values of the relative value loss $\mathbf{L}_\gamma(\mathcal{F}, \mathcal{F}_G)$ as a function of the number of providers, for instances with 30 jobs constructed from the data, using a time window of $w = 16$ minutes. **(right)** Average value of the relative value loss $\mathbf{L}_\gamma(\mathcal{F}, \mathcal{F}_G)$ as a function of the number of providers for synthetic instances with different coefficients of variation (cv), and using the parametric value $x = 1.6$.

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 Avenues for Future Research

In this paper we established that the relative value loss due to a broad class of provider guarantees is bounded. We further showed that the worst-case losses are primarily driven by fairness considerations when a many providers are present, and by the heterogeneity in the providers' ability 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 carry a nonintuitive impact, leading to *smaller*

losses. Finally, we confirmed several of the findings numerically using both real and synthetic data, wherein we also documented that the value loss in a specific problem setting may be significantly lower than the worst-case values we obtained.

These results motivate future work from both a theoretical and a practical perspective. Having established that losses are bounded in a general setting and under a broad class of provider guarantees, one could now aim to understand how the losses behave in more particular settings or for subclasses of guarantees obtained as special cases of our framework. For instance, one could seek to establish a parametrized upper-bound on the relative loss when guaranteeing a specific 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 §5. Additionally, and from a more prescriptive viewpoint, losses being low for many types of guarantees also opens 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.

References

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

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

Awerbuch B, Azar Y, Richter Y, Tsur D (2006) Tradeoffs in worst-case equilibria. *Theoretical Computer Science* 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. *Journal of Organizational Behavior* 12(1):39–53.

Berger PD (1972) On setting optimal sales commissions. *Journal of the Operational Research Society* 23(2):213–215.

Bertsekas D, Nedić A, Ozdaglar A (2003) *Convex Analysis and Optimization.* Athena Scientific optimization and computation series (Athena Scientific).

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

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

Besbes O, Iancu DA, Trichakis N (2017) Dynamic pricing under debt: Spiraling distortions and efficiency losses. *Forthcoming in Management Science* .

Brockner J, Wiesenfeld BM (1996) An integrative framework for explaining reactions to decisions: interactive effects of outcomes and procedures. *Psychological bulletin* 120(2):189.

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

Chawla S, Roughgarden T (2008) Bertrand competition in networks. *International Symposium on Algorithmic Game Theory*, 70–82 (Springer).

CNBC (2017) Only 4% of Uber drivers remain on the platform a year later, says report. `https://www.cnbc.com/2017/04/20/only-4-percent-of-uber-drivers-remain-after-a-year-says-report.html`, [Web page accessed 7-Jul-2018].

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

Correa JR, Schulz AS, Stier-Moses NE (2004) Selfish routing in capacitated networks. *Mathematics of Operations Research* 29(4):961–976.

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

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

Financial Times (2014) Uber offers US drivers perks to boost retention. `https://www.ft.com/content/446dd780-6f91-11e4-8d86-00144feabdc0`, [Web page accessed 7-Jul-2018].

Gibbons R (1998) Incentives in organizations. *The Journal of Economic Perspectives* 12(4):115–132.

Golovin D (2005) Max-min fair allocation of indivisible goods. Technical Report Research Showcase at CMU, Carnegie Mellon University, Pittsburgh, PA.

Greenberg J (1982) Chapter 11 - Approaching Equity and Avoiding Inequity in Groups and Organizations. Greenberg J, Cohen RL, eds., *Equity and Justice in Social Behavior*, 389 – 435 (Academic Press).

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

Johari R, Tsitsiklis JN (2004) Efficiency loss in a network resource allocation game. *Mathematics of Operations Research* 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. *Computer Science Review* 3(2):65 – 69.

Lazear EP (2000) Performance pay and productivity. *American Economic Review* 90(5):1346–1361.

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

Lipton RJ, Markakis E, Mossel E, Saberi A (2004) On approximately fair allocations of indivisible goods. *Proceedings of the 5th ACM conference on Electronic commerce - EC '04* (Aamas):125.

Lyft Inc (2018) Weekly Guarantees. `https://help.lyft.com/hc/en-ca/articles/115012927247-Weekly-Guarantees`, [Web page accessed 7-Jul-2018].

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

NYC Taxi and Limousine Commission (2016) TLC Trip Record Data. `http://www.nyc.gov/html/tlc/html/about/trip_record_data.shtml`, [Database retrieved on 10-November-2017].

Papadimitriou C (2001) Algorithms, games, and the internet. *Proceedings of the Thirty-third Annual ACM Symposium on Theory of Computing*, 749–753, STOC '01 (New York, NY, USA: ACM).

Patagonia (2018) Corporate Responsibility. `https://www.patagonia.com/corporate-responsibility.html`, [Web page accessed 27-Jul-2018].

Perakis G, Roels G (2007) The price of anarchy in supply chains: Quantifying the efficiency of price-only contracts. *Management Science* 53(8):1249–1268.

Rawls J (1971) *A Theory of Justice*. Oxford Paperbacks 301 301 (Harvard University Press).

Rideshare Central (2018) Uber Sign Up Bonus. `https://ridesharecentral.com/uber-sign-up-bonus`, [Web page accessed 25-Feb-2019].

Roughgarden T (2003) The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences* 67(2):341 – 364, special Issue on STOC 2002.

Roughgarden T, Tardos E (2002) How bad is selfish routing? *Journal of ACM* 49(2):236–259.

Simon HA (1956) Rational choice and the structure of the environment. *Psychological Review* 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. *Journal of Occupational and Organizational Psychology* 70(4):391–408.

Starbucks Corporation (2018) Starbucks Ethical Sourcing of Sustainable Products. `https://www.starbucks.com/responsibility/sourcing`, [Web page accessed 27-Jul-2018].

The Rideshare Guy (2016) Why Do So Many People Quit Driving For Uber? `https://therideshareguy.com/why-do-so-many-people-quit-driving-for-uber/`, [Web page accessed 25-Feb-2019].

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 & Organization Management* 25(3):269–290.

Uber Technologies Inc (2018a) Driver-Partner Incentives. `https://www.uber.com/drive/atlanta/resources/driver-partner-incentive-guarantee-faq-questions/`, [Web page accessed 7-Sept-2018].

Uber Technologies Inc (2018b) Hourly Guarantees. `https://www.uber.com/en-MY/drive/resources/earn-with-hourly-guarantees/`, [Web page accessed 24-May-2018].

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 Operations 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.

Appendices

A   Proofs and Examples

For ease of notation, let $\mathcal{F}^* \stackrel{\text{def}}{=} \arg\max\limits_{\boldsymbol{A} \in \mathcal{F}} \sum\limits_{i=1}^{n} \gamma_i v(A_i)$, and $\mathcal{F}_{\text{G}}^* \stackrel{\text{def}}{=} \arg\max\limits_{\boldsymbol{B} \in \mathcal{F}_{\text{G}}} \sum\limits_{i=1}^{n} \gamma_i v(B_i)$.

We begin by proving an optimality condition that any allocations $\boldsymbol{A} \in \mathcal{F}^*$ and $\boldsymbol{B} \in \mathcal{F}_{\text{G}}^*$ must satisfy.

Lemma 1 *For any fixed $\mathcal{F}$ and $\mathcal{F}_{\text{G}}$, let $\boldsymbol{A} \in \mathcal{F}^*$, and $\boldsymbol{B} \in \mathcal{F}_{\text{G}}^*$, then*

$$v(A_i \setminus (B_1 \cup \ldots \cup B_n)) \leq v(B_j \setminus A_i), \quad \forall\, i, j \in N \tag{10}$$

Proof. Assume by contradiction that $v(A_i \setminus (B_1 \cup \ldots \cup B_n)) > v(B_j \setminus A_i)$, for some i, j. Thus, take $B_j' = (B_j \cap A_i) \cup (A_i \setminus (B_1 \cup \ldots \cup B_n))$. Then, $v(B_j') = v((B_j \cap A_i)) + v((A_i \setminus (B_1 \cup \ldots \cup B_n))) > v((B_j \cap a_i)) + v((B_j \setminus A_i)) = v(B_j)$. Additionally, $B_j' \subseteq A_i$, implying by Assumption 1-(iii) that (B_j', A_{-j}) is 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 $(B_1, \ldots, B_j', \ldots, B_n) \in \mathcal{F}$, where B_j is replaced by B_j'. But then, by Assumption 2 on $\mathcal{F}_{\text{G}}$, we must have that $(B_1, \ldots, B_j', \ldots, B_n) \in \mathcal{F}_{\text{G}}$. This implies a contradiction, because $(B_1, \ldots, B_n) \in \mathcal{F}_{\text{G}}^*$, but $\sum_{i=1}^{n} \gamma_i v(B_i) < \sum_{i=1}^{j-1} \gamma_i v(B_i) + \gamma_j v(B_j') + \sum_{i=j+1}^{n} \gamma_i v(B_i)$. ∎

Using Lemma 1, we now prove Theorem 1, that shows an upper bound of $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\text{G}})$.

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

$$x_i = v(A_i \setminus \cup_{k=1}^{n} B_k), \text{ for } i \in N$$

$$y_i = v(B_i \setminus \cup_{k=1}^{n} A_k), \text{ for } i \in N$$

$$w_{ij} = v(A_i \cap B_j) \text{ for } i, j \in N$$

With these variables, we can rewrite the inequalities proven in Lemma 1 as:

$$x_i + w_{ij} - \sum_{k=1}^{n} w_{kj} - y_j \leq 0, \text{ for } i, j \in N. \tag{11}$$

Moreover, we can write the expression in (1) that defines $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\text{G}})$ using these same variables

as:

$$\frac{\displaystyle\sum_{i=1}^{n}\gamma_i x_i - \sum_{i=1}^{n}\gamma_i y_i + \sum_{i=1}^{n}\sum_{j=1}^{n} w_{ij}(\gamma_i - \gamma_j)}{\displaystyle\sum_{i=1}^{n}\gamma_i x_i + \sum_{i=1}^{n}\sum_{j=1}^{n} w_{ij}\gamma_i} \tag{12}$$

Notice that written in this way it is clear that $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}})$ is a Fractional Linear function of the variable x, y, and w. Because for any $\mathcal{F}$ and $\mathcal{F}_{\mathrm{G}}$, the inequalities of (11) must hold, then, we can find an upper bound on $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}})$ for any $\mathcal{F}$ and $\mathcal{F}_{\mathrm{G}}$, given γ, by solving the following Fractional Linear Program:

$$\begin{aligned}
\text{Maximize} \quad & \frac{\displaystyle\sum_{i=1}^{n}\gamma_i x_i - \sum_{i=1}^{n}\gamma_i y_i + \sum_{i=1}^{n}\sum_{j=1}^{n} w_{ij}(\gamma_i - \gamma_j)}{\displaystyle\sum_{i=1}^{n}\gamma_i x_i + \sum_{i=1}^{n}\sum_{j=1}^{n} w_{ij}\gamma_i} \\
\text{subj. to} \quad & x_i + w_{ij} - \sum_{k=1}^{n} w_{kj} - y_j \leq 0 && \text{for } i,j \in N \\
& x_i \geq 0 && \text{for } i \in N \\
& y_i \geq 0 && \text{for } i \in N \\
& w_{ij} \geq 0 && \text{for } i,j \in N
\end{aligned} \tag{13}$$

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

$$\begin{aligned}
\text{Maximize} \quad & \sum_{i=1}^{n}\gamma_i x_i - \sum_{i=1}^{n}\gamma_i y_i + \sum_{i=1}^{n}\sum_{j=1}^{n} w_{ij}(\gamma_i - \gamma_j) \\
\text{subj. to} \quad & \sum_{i=1}^{n}\gamma_i x_i + \sum_{i=1}^{n}\sum_{j=1}^{n} w_{ij}\gamma_i = 1 \\
& x_i + w_{ij} - \sum_{k=1}^{n} w_{kj} - y_j \leq 0 && \text{for } i,j \in N \\
& x_i \geq 0 && \text{for } i \in N \\
& y_i \geq 0 && \text{for } i \in N \\
& w_{ij} \geq 0 && \text{for } i,j \in N.
\end{aligned} \tag{14}$$

The optimization problem in (14) is a linear program with variables x, y, and w, and an objective function bounded above by 1. Thus, we know that strong duality must hold and we can find the

desired upper bound by studying the dual linear program:

$$
\begin{aligned}
\text{Minimize} \quad & s \\
\text{subj. to} \quad & s\gamma_i + \sum_{j=1}^{n} \lambda_{ij} \geq \gamma_i && \text{for } i \in N \\
& \gamma_j - \gamma_k + s\gamma_k - \sum_{\substack{i=1 \\ i \neq k}}^{n} \lambda_{ij} \geq 0 && \text{for } k, j \in N \\
& \gamma_i \geq \sum_{j=1}^{n} \lambda_{ij} && \text{for } i \in N \\
& \lambda_{ij} \geq 0 && \text{for } i, j \in N.
\end{aligned}
\tag{15}
$$

Problem (15) in turn can be rewritten as:

$$
\begin{aligned}
\text{Minimize} \quad & \max\left\{ \left\{ 1 - \frac{1}{\gamma_i} \sum_{j=1}^{n} \lambda_{ij} \right\}_{i=1}^{n}, \left\{ \frac{1}{\gamma_k} \left(\sum_{\substack{i=1 \\ i \neq k}}^{n} \lambda_{ij} + \gamma_k - \gamma_j \right) \right\}_{k,j=1}^{n} \right\} \\
\text{subj. to} \quad & \gamma_i \geq \sum_{j=1}^{n} \lambda_{ij} && \text{for } i \in N \\
& \lambda_{ij} \geq 0 && \text{for } i, j \in N
\end{aligned}
\tag{16}
$$

Let us define $f(\boldsymbol{\gamma}) : [\gamma_{\min}, \gamma_{\max}]^n \to [0, 1]$, as the solution to the dual problem (16). Now we can see that each term that appears in the objective of problem (16) is a quasiconvex function of both $\boldsymbol{\gamma}$ 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 (16) is convex in $\boldsymbol{\gamma}$ and λ, which means that the function $f(\boldsymbol{\gamma})$ must be quasiconvex in $\boldsymbol{\gamma}$, because it is the minimization of a quasiconvex function over a convex set. But then, if we wish to find $\max_{\boldsymbol{\gamma} \in [\gamma_{\min}, \gamma_{\max}]^n} f(\boldsymbol{\gamma})$, 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 2, if we take the instance where the first n_0 values of $\boldsymbol{\gamma}$ are $\gamma_{\max}$, and the rest are $\gamma_{\min}$, for $n_0 \in N$, we can see that the optimal value is $\max\left\{ \delta, \frac{n-1}{n+n_0+(1-\delta)(n-n_0)-1} \right\}$. But, notice that this is a decreasing function of n_0, implying that the instance of $\boldsymbol{\gamma}$ that maximizes the solution to problem (13) is when $n_0 = 1$. In this case, we recover $\max\left\{ \delta, \frac{n-1}{n+(1-\delta)(n-1)} \right\}$. Which proves the theorem. ∎

Lemma 2 *The optimal value of the linear program (14), when $\gamma_i = \gamma_{\max}$ for all $i \in \{1, \ldots, n_0\}$,*

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 (15) is a dual of the linear program (14). 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:

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

By simply replacing this values in problem (14), 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:

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

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

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

Finally, we take the following dual solution:

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

35

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

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

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. ∎

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 $\mathcal{F}$, $\mathcal{F}_{\mathrm{G}}$, and $\boldsymbol{\gamma}$ such that $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}})$ achieves the values in (6)-(7). For this, we will generalize Instances 1 and 2, for n agents.

We begin by considering an instance family that achieves $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) = \delta$, for any n and δ.

Instance 6 *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 $i \in \{2,\ldots,n\}$. Given this D with only one job, consider $\mathcal{F} = \{\boldsymbol{A}, \boldsymbol{B}\}$, where $A_1 = \{d_1\}$, $A_i = \emptyset$ for all $i \in \{2,\ldots,n\}$, and $B_1 = \emptyset$, $B_2 = \{d_1\}$, and $B_i = \emptyset$ for all $i \in \{3,\ldots,n\}$. Finally, if we take $\mathcal{F}_{\mathrm{G}} = \{\boldsymbol{B}\}$, then the only efficient allocations would be $\boldsymbol{A}$, that gives a value of $\gamma_{\max}$, while, by definition, the only allocation in $\mathcal{F}_{\mathrm{G}}$ would be $\boldsymbol{B}$, which implies that $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) = \frac{\gamma_{\max} - \gamma_{\min}}{\gamma_{\max}} = \delta$.*

Now we present a family of instances that have $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) = \frac{n-1}{n+(1-\delta)(n-1)} - \epsilon$, for any $\epsilon > 0$, where $\mathcal{F}_{\mathrm{G}} = \mathcal{F}_{\mathrm{G}}^{\mathrm{mM}}$ 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, while 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 7 *Given any n, let $D = \{d_1,\ldots,d_n,\ldots,d_{2n-1}\}$, such that $v(d_i) = 1$, for all $i \in N$, and $v(d_j) = 1 - \kappa$, for $j \in \{n+1,\ldots,n\}$. Let as well $\gamma_1 = \gamma_{\max} \geq \gamma_{\min} = \gamma_i$, for all $i \in \{2,\ldots,n\}$. A subset $A \subseteq D$ can be assigned to a single provider if either $d_j \notin A$, for all $j \in \{n+1,\ldots,2n-1\}$, or $A = \{d_i\}$ for some $i \in \{n+1,\ldots,2n-1\}$. Let $\mathcal{F}$ be formed by all possible disjoint combinations of*

such subsets of D. Let $p : \mathcal{P}(D) \to \mathbb{R}$ 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 $\mathcal{F}_\mathrm{G}^\mathrm{mM}(p, N)$ be the associated Max-Min fair restriction, as described in (4).

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

Instance 8 Given any n, and t, positive integers, let $D = \cup_{k=1}^t C^k \cup S$, where $C^k = \{d_2^k, \ldots, d_n^k\}$, and $S = \{d_1^s, \ldots, d_{n(t+1)}^s\}$. Let as well $\gamma_1 = \gamma_\mathrm{max} \geq \gamma_\mathrm{min} = \gamma_i$, for all $i \in \{2, \ldots, n\}$. A subset of jobs $A \subset D$ 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 $k \in \{1, \ldots, t\}$ 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 $p : \mathcal{P}(D) \to \mathbb{R}$ 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 $\mathcal{F}_\mathrm{G}^\mathrm{mM}(p, N)$ be the associated Max-Min fair restriction, as described in (4).

Given this instance, the only efficient allocation (modulo symmetries) would be $\boldsymbol{A}$, such that $A_1 = S$, $A_i = \{d_i^1, d_i^2, \ldots, d_i^t\}$, for all $i \in \{2, \ldots, n\}$. 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 $\mathcal{F}_\mathrm{G}^\mathrm{mM}$ are of the form $\boldsymbol{B}$, 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, while in any Max-Min fair allocation, all providers are being allocated exactly $t+1$ jobs. Hence, the loss generated by this instance, that comes from the fact that none of the jobs in $\cup_{k=1}^t C^k$ are allocated for any allocation in $\mathcal{F}_\mathrm{G}^\mathrm{mM}$, is $\frac{(t+1)n\gamma_\mathrm{max} + t(n-1)\gamma_\mathrm{min} - (t+1)(\gamma_\mathrm{max} + (n-1)\gamma_\mathrm{min})}{(t+1)n\gamma_\mathrm{max} + t(n-1)\gamma_\mathrm{min}} = \frac{t(n-1) + \delta(n-1)}{t(n+(n-1)(1-\delta)) + n}$.

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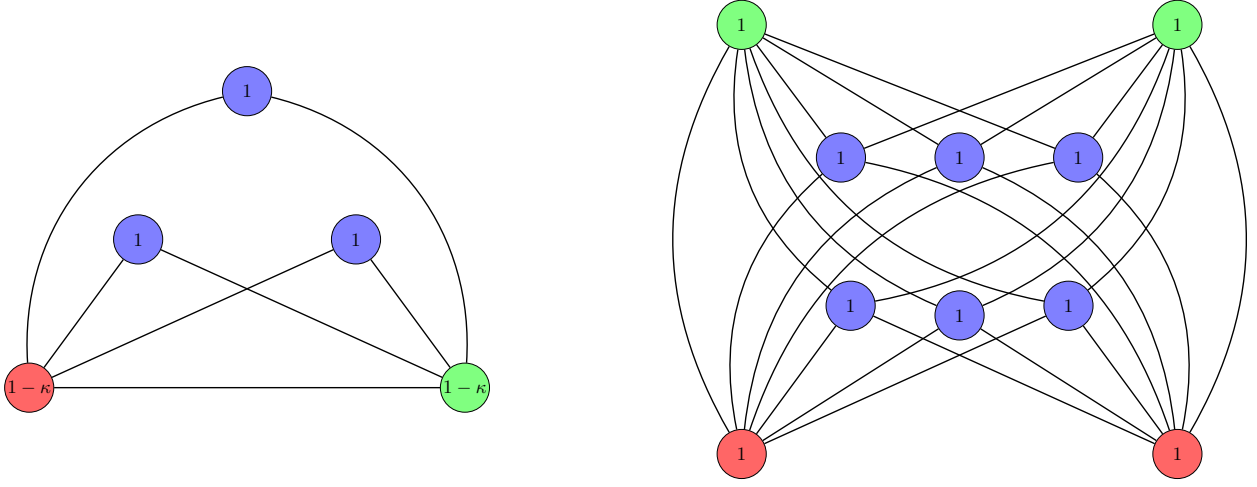

Figure 9: **Feasibility graphs for Instance 7 (left)**, when n is 3, and Instance 8 **(right)**, when n is 3 and t is 1.

To better visualize the structure of Instances 7 and 8, 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 7, there would be no edges between any pair of the first n vertices, while 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 3 providers of this graph can be seen in Figure 9. These types of graphs are known as complete split graphs (see Le and Peng (2015)). Similarly, in the case of Instance 8, there would be no edges between any pair of vertices in the set S, while 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 3 providers, and $t = 1$ can be seen in Figure 9. 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 6, 7, and 8 prove that $\mathbf{L}_\gamma(\mathcal{F}, \mathcal{F}_G)$ can be taken as close to $\max\left\{\delta, \frac{n-1}{n+(1-\delta)(n-1)}\right\}$ as desired. This concludes the proof of Theorem 2. ∎

We now proof Proposition 1.

Proof of Proposition 1. If we have that $\mathcal{F}_G = \arg\max\limits_{\boldsymbol{A} \in \mathcal{F}} g(\boldsymbol{A})$, then if $\boldsymbol{A} \in \mathcal{F}_G$, and $\boldsymbol{B} \in \mathcal{F}$ satisfy that $v(B_i) \geq v(A_i)$, for each i, then we know by the statement of the Proposition that $g(\boldsymbol{B}) \geq g(\boldsymbol{A})$,

which implies that $\boldsymbol{B} \in \arg\max\limits_{\boldsymbol{A} \in \mathcal{F}} g(\boldsymbol{A}) = \mathcal{F}_\mathrm{G}$, proving that Assumption 2 holds. On the other hand, if we assume that Assumption 2 holds for a certain $\mathcal{F}_\mathrm{G}$, then let us consider the function $g : \mathcal{F} \to \mathbb{R}$, defined by $g(\boldsymbol{A}) = \mathbb{1}_{\mathcal{F}_r}(\boldsymbol{A})$, that takes the value 1 when $\boldsymbol{A} \in \mathcal{F}_\mathrm{G}$, and 0 otherwise. Thus, by definition $\mathcal{F}_\mathrm{G} = \arg\max\limits_{\boldsymbol{A} \in \mathcal{F}} g(\boldsymbol{A})$. Moreover, if $g(\boldsymbol{A}) = 1$, and $\boldsymbol{B} \in \mathcal{F}$ is such that $v(B_i) \geq v(A_i)$, for all $1 \leq i \leq n$, then, because $\mathcal{F}_\mathrm{G}$ satisfies Assumption 2, $\boldsymbol{B} \in \mathcal{F}_\mathrm{G}$, which implies that $g(\boldsymbol{B}) = 1 \geq g(\boldsymbol{A})$. Therefore, $g()$ satisfies the condition of Proposition 1, which concludes the proof. ∎

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 $\mathcal{F}_\mathrm{G}$ that satisfies this assumption can be expressed as such union. For this, take any $\mathcal{F}_\mathrm{G}$ that satisfies Assumption 2, and consider for each $\boldsymbol{A} \in \mathcal{F}_\mathrm{G}$, the guarantee $\mathcal{F}_\mathrm{G}{}^{\boldsymbol{A}} = \{\boldsymbol{B} \in \mathcal{F} \mid v(B_i) \geq v(A_i), \forall i \in N\}$. We claim than $\cup_{\boldsymbol{A} \in \mathcal{F}_\mathrm{G}} \mathcal{F}_\mathrm{G}{}^{\boldsymbol{A}} = \mathcal{F}_\mathrm{G}$. Clearly $\mathcal{F}_\mathrm{G} \subseteq \cup_{\boldsymbol{A} \in \mathcal{F}_\mathrm{G}} \mathcal{F}_\mathrm{G}{}^{\boldsymbol{A}}$, because each $\boldsymbol{A} \in \mathcal{F}_\mathrm{G}{}^{\boldsymbol{A}}$. Moreover, if $\boldsymbol{C} \in \cup_{\boldsymbol{A} \in \mathcal{F}_\mathrm{G}} \mathcal{F}_\mathrm{G}{}^{\boldsymbol{A}}$, then there exists $\boldsymbol{A} \in \mathcal{F}_\mathrm{G}$, such that $v(C_i) \geq v(A_i), \forall i \in N$, but then by Assumption 2, $\boldsymbol{C} \in \mathcal{F}_\mathrm{G}$, which concludes the proof. ∎

We now present Instance 9, that shows how when we relax Assumption 1 (ii), we can achieve a loss that grows asymptotically to 100% with the number of providers.

Instance 9 *Consider $n = 2k - 1$, for any integer $k > 0$, $D = \{d_1^1, \ldots, d_k^1, \ldots, d_1^k, \ldots, d_k^k\}$, and*

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

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

$$\mathbf{L}_\gamma(\mathcal{F}, \mathcal{F}_\mathrm{G}^{\mathrm{mM}}(p, N)) = 1 - \frac{4n}{(n+1)^2}.$$

Where $\mathcal{F}_\mathrm{G}^{\mathrm{mM}}$ 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) $\boldsymbol{B}$, such that $B_i = \{d_k^i\}$, for $i \in \{1, \ldots, k\}$, and $B_j = \{d_{j-k}^1, \ldots, d_{j-k}^k\}$, for $j \in \{j+1, \ldots, 2k-1\}$. 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 $\boldsymbol{A}$ such that $A_i = \{d_1^i, \ldots, d_k^i\}$, for $i \in \{1, \ldots, k\}$. This would lead to a total value generated of exactly k, which would imply that

$$\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}^{\mathrm{mM}}(p, N)) = \frac{k - \frac{2k-1}{k}}{k}$$

$$= 1 - \frac{4n}{(n+1)^2}.$$

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

We will now prove Propositions 3 to 7 from §4.

Proof of Proposition 3. In order to prove this proposition, we will first formally define the set of guarantees $\mathcal{F}_{\mathrm{G}}^{\mathrm{c}}$, as any subset of $\mathcal{F}^{\mathrm{c}}$, that satisfies Assumption 2, replacing in the definition of the assumption $\mathcal{F}$ by $\mathcal{F}^{\mathrm{c}}$, and using the generalized notion of $v(C_i)$ for $C \in \mathcal{F}^{\mathrm{c}}$, described in §4. In particular, we extend the notion of uniform income guarantees under monotonic payment functions: given $p(\cdot)$, a monotonic payment function, we take $\mathcal{F}_{\mathrm{G}}^{\mathrm{c}} = \{C \in \mathcal{F}^{\mathrm{c}} \mid p(C_i) \geq \tau, \text{ for } i \in N\}$. As in §2.3, 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, $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}^{\mathrm{c}}, \mathcal{F}_{\mathrm{G}}^{\mathrm{c}}) = 0$.

Consider any allocation $\boldsymbol{A} \in \mathcal{F}^*$. We know that any permutation $\boldsymbol{A}_\sigma$ of $\boldsymbol{A}$ is as well in $\mathcal{F}$, hence, take $\boldsymbol{C} \in \mathcal{F}^{\mathrm{c}}$, such that $\boldsymbol{C} = (\{\theta_\sigma\}_{\sigma \in S_n}, \{A_\sigma\}_{\sigma \in S_n})$, 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 $\boldsymbol{C}$ must be in any non empty uniform income guarantee. To see this, let us assume by contradiction that there is a nonempty $\mathcal{F}_{\mathrm{G}}^{\mathrm{c}}$ such that $\boldsymbol{C} \notin \mathcal{F}_{\mathrm{G}}^{\mathrm{c}}$. 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, $\boldsymbol{C} \notin \mathcal{F}_{\mathrm{G}}^{\mathrm{c}}$ 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 $\boldsymbol{B} \in \mathcal{F}^{\mathrm{c}}$ 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 $\boldsymbol{B}^j \in \mathcal{F}$, that achieves a higher total value than $\boldsymbol{A} \in \mathcal{F}^*$, which leads to a contradiction and proves the proposition. ∎

Proof of Proposition 4.

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

To see this notice that the total amount of jobs that can be allocated is $\hat{k}$, and thus, if the total value generated in any $\boldsymbol{A} \in \mathcal{F}^*$ were 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 $\{d_1, \ldots, d_{\hat{k}}\} \setminus (\cup_{i=1}^{n} A_i)$. This replacement would generate a feasible allocation, and would improve the total value generated, which leads to a contradiction because $\boldsymbol{A} \in \mathcal{F}^*$.

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

Take any allocation $\boldsymbol{B} \in A$, if $(\cup_{i=1}^{n} B_i) \setminus \{d_1, \ldots, d_{\hat{k}}\} = \emptyset$, then $\boldsymbol{B}$ allocates only elements of $\{d_1, \ldots, d_{\hat{k}}\}$. Otherwise, consider $\boldsymbol{A}$ such that we replace in $\boldsymbol{B}$ every job in $(\cup_{i=1}^{n} B_i) \setminus \{d_1, \ldots, d_{\hat{k}}\}$ by an element in $\{d_1, \ldots, d_{\hat{k}}\} \setminus (\cup_{i=1}^{n} A_i)$. 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 $\boldsymbol{A} \in \mathcal{F}$, such that only jobs in $\{d_1, \ldots, d_{\hat{k}}\}$ are allocated, then if there are any jobs in the first $\hat{k}$ not allocated in $\boldsymbol{C}$, this means that there is at least one provider i such that $|A_i| < k_i$. Consider then the allocation $\boldsymbol{C} \in \mathcal{F}$, such that we add jobs from $\{d_1, \ldots, d_{\hat{k}}\}$ to $\boldsymbol{A}$, until all $|A_i| = k_i$. This allocation $\boldsymbol{C}$ Pareto dominates allocation $\boldsymbol{A}$, and uses exactly all elements in $\{d_1, \ldots, d_{\hat{k}}\}$. Hence, as mentioned above, this proves that $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) = 0$, for any $\mathcal{F}_{\mathrm{G}}$, satisfying Assumption 2. ∎

Proof of Proposition 5.

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 $\mathcal{F}_{\mathrm{G}}$, we will show that when $|D| < n$, $\max_{\boldsymbol{A} \in \mathcal{F}} \min_{i=1}^{n} v(A_i) = 0$, and when $|D| \geq n$, $\max_{\boldsymbol{A} \in \mathcal{F}} \min_{i=1}^{n} v(A_i) = 1$.

When $|D| < n$, any allocation $\boldsymbol{A} \in \mathcal{F}$ will necessarily have an $A_i = \emptyset$, which implies that

$\max_{\boldsymbol{A}\in\mathcal{F}}\min_{i=1}^{n} v(A_i) = 0$. On the other hand, if $|D| \geq n$, then for any $\boldsymbol{A} \in \mathcal{F}$, 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, $\max_{\boldsymbol{A}\in\mathcal{F}}\min_{i=1}^{n} v(A_i) = 1$.

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

Now, to show (ii), we simply observe that Instance 5 is a generalized form of Instance 2, and that by taking $v(d) = 1 - \kappa$, for all $d \in C$, $v(d) = 1$, for all $d \in S$, and $|C| = n - 1$, $|S| = n$, we obtain exactly Instance 2, when $n = 2$, and Instance 7 when $n \geq 2$. ∎

We now describe Example 2, that shows how the dependency of the loss on the variance of the values described in Example 1 extends to the n provider case.

Example 2 *Consider a variant of Instance 7, 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 $i \in \{n+1, \ldots, 2n-1\}$. Hence, if we take the expectation of the loss, with respect to the error κ, we get*

$$
\begin{aligned}
\mathbb{E}_{\kappa}(\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}})) &= \int_{0}^{\min\{1, \frac{\Delta}{2}\}} \frac{1}{\Delta} \frac{(n-1)(1-\kappa)}{n + (n-1)(1-\kappa)} \, d\kappa \\
&= \frac{1}{\Delta}\left(\frac{n}{n-1}\log((n-1)\min\{1, \frac{\Delta}{2}\} - (2n-1)) + \min\{1, \frac{\Delta}{2}\} - \frac{n}{n-1}\log(2n-1)\right) \\
&:= g(\Delta).
\end{aligned}
$$

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 values, for higher variance, there is a lower expected loss.

Proof of Proposition 6. We will first prove statement (i), namely, that if $n = 1$, and $\gamma_{m}in = \gamma_{m}ax$, then $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) = 0$, for any symmetric set of allocations with guarantees, $\mathcal{F}_{\mathrm{G}}$. To show this, we simply observe that due to Assumption 2, any $\boldsymbol{A} \in \mathcal{F}^*$, must also satisfy $\boldsymbol{A} \in \mathcal{F}_{\mathrm{G}}$, because clearly for any $\boldsymbol{B} \in \mathcal{F}_{\mathrm{G}}$, $v(B_1) \leq v(A_1)$.

Now, we will show that (ii) holds. For this, 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 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 $\max_{\boldsymbol{A} \in \mathcal{F}} \sum_{i} v(A_i) = \sum_{d \in D} v(d)$. 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 $\mathcal{F}_{\mathrm{G}}^*$ there exists an allocation $\boldsymbol{B}$, 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 $\boldsymbol{A} \in \mathcal{F}_{\mathrm{G}}^*$, 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 $\boldsymbol{A}'$ such that $A'_j = A_j$, for $j \not= i$, and $A'_i = \{d\}$. This leads to a contradiction, because by Assumption 2, $\boldsymbol{A}' \in \mathcal{F}_{\mathrm{G}}$, and $\sum_{i \in N} v(A'_i) > \sum_{i \in N} v(A_i)$, but $\boldsymbol{A} \in \mathcal{F}_{\mathrm{G}}^*$. Therefore, there exists an allocation $\boldsymbol{B} \in \mathcal{F}_{\mathrm{G}}^*$, 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 $\mathbf{L}_{\gamma}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) = 0$, for any set of allocations with guarantees $\mathcal{F}_{\mathrm{G}}$.

∎

Proof of Proposition 7. 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 $m < 2n$, and because we can always allocate only one job to any provider, then $\max_{\boldsymbol{A} \in \mathcal{F}} \min_{i \in N} v(A_i) = 1$. 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 $\mathbf{L}_{\gamma}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) = 0$, under any uniform income-guarantee under monotonic payment functions. ∎

B Numerical Analysis of Synthetic and Real-world Data

In this section we provide the details of the numerical analysis we discuss in §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.

Instances generated with real-world data. We used the publicly available dataset provided by NYC Taxis and Limousine Commission. This dataset includes, for each yellow-taxi ride, the total fare, the starting and ending location 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 am and 5 pm, in order to restrict our attention to a time horizon with a relatively constant rate of trips per time. Moreover, for each date, 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 Assumption 1(ii), that any feasible set of jobs could be performed by any provider. In Figure 10 we can see the empirical distribution of the total trip duration and total fare payed, 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, while 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, this removed several outliers that were clearly due to corrupted data (trips of almost 24 hours or more than $1000).
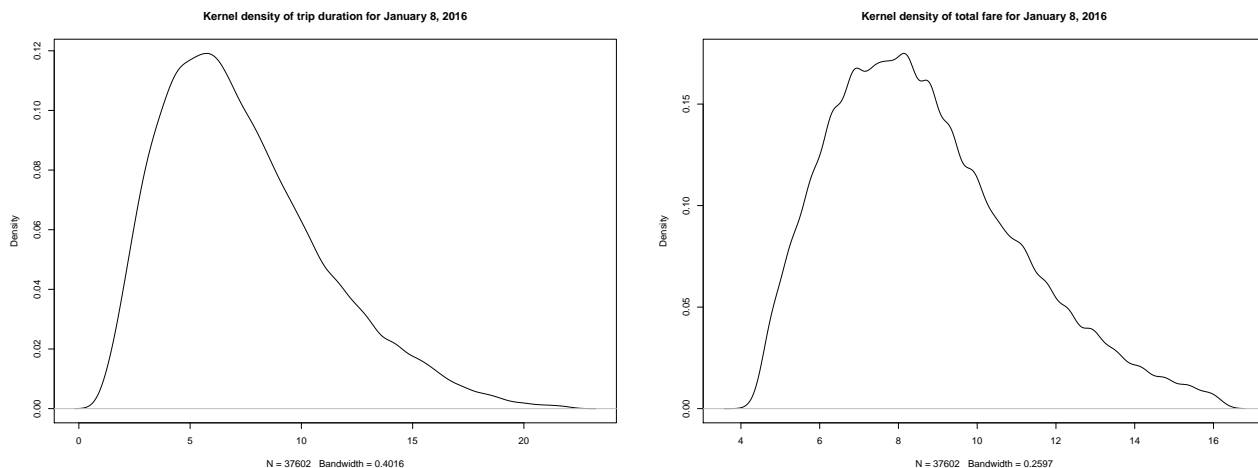
Figure 10: **(left)** Empirical distribution of the duration of trips. **(right)** Empirical distribution of the total fare trips.

In order to compute the average relative loss when considering only Max-Min fair solutions, we generated demand instances using this 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 to satisfy that no trips that intersected in time would be allocated to the same provider. 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 a 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). Once 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 11-12 we can see the relative value loss as a function of the number of providers, for different combinations of dates, regions, and size of interval, w. We show here a representative set of our results, the complete set of results is available upon request.

We can see in Figure 11 that there doesn't 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 12 we see that for the same region but different dates the relative losses do not appear to change much. Nevertheless, by comparing Figure 11 to Figure 12 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 12 and 11 we observed the same pattern mentioned in §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 it's 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 due to frequent low losses, as opposed to infrequent high losses.
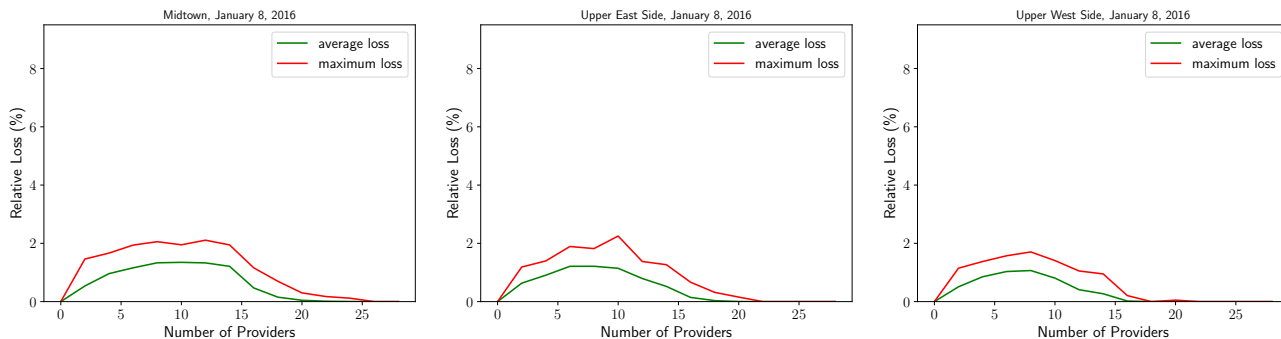

Figure 11: **Relative value loss for different regions.** Average and maximum $\mathbf{L}_\gamma(\mathcal{F}, \mathcal{F}_G)$ as a function of the number of providers, for instances with 30 jobs constructed from the data, using $w = 20$ and three different regions of NYC.

Synthetically generated Instances. In order to analyze the dependency between the variation of values and the relative value loss we mention in §4, we generate synthetic instances where we can
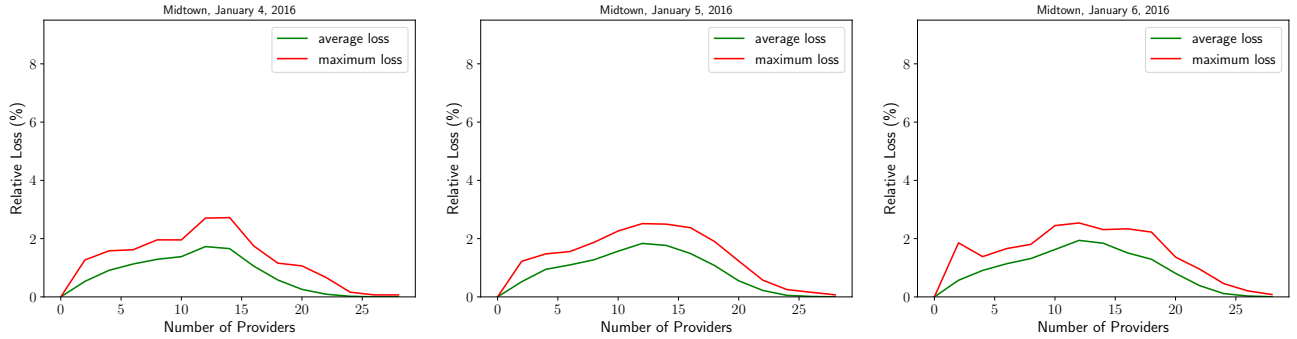
Figure 12: **Relative value loss for different dates.** Average and maximum $\mathbf{L}_\gamma(\mathcal{F}, \mathcal{F}_G)$ as a function of the number of providers, for instances with 30 jobs constructed from the data, using $w = 15$ and three different dates of the first week of 2016.

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 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 13. 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 due to the fact that we take instances of 30 jobs for all the values of x, which implies that the probability that two jobs are incompatible is lower for larger x.

We can see in Figure 13, as we mentioned in §5, that the maximum loss in decreasing in the coefficient of variation of the values, consistent with the remarks of §4 on the variation of values as a driver of loss. Nevertheless, in the instances we generated for Figure 13, the variation of the values is intrinsically connected to the variation in the sizes 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 1, for feasibility purposes. We then plotted the average losses as before in Figure 14. By comparing

46

Figure 13 and 14, 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 1. 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. And moreover it demonstrates the effect of many of the main drivers of loss we analyzed in §4.
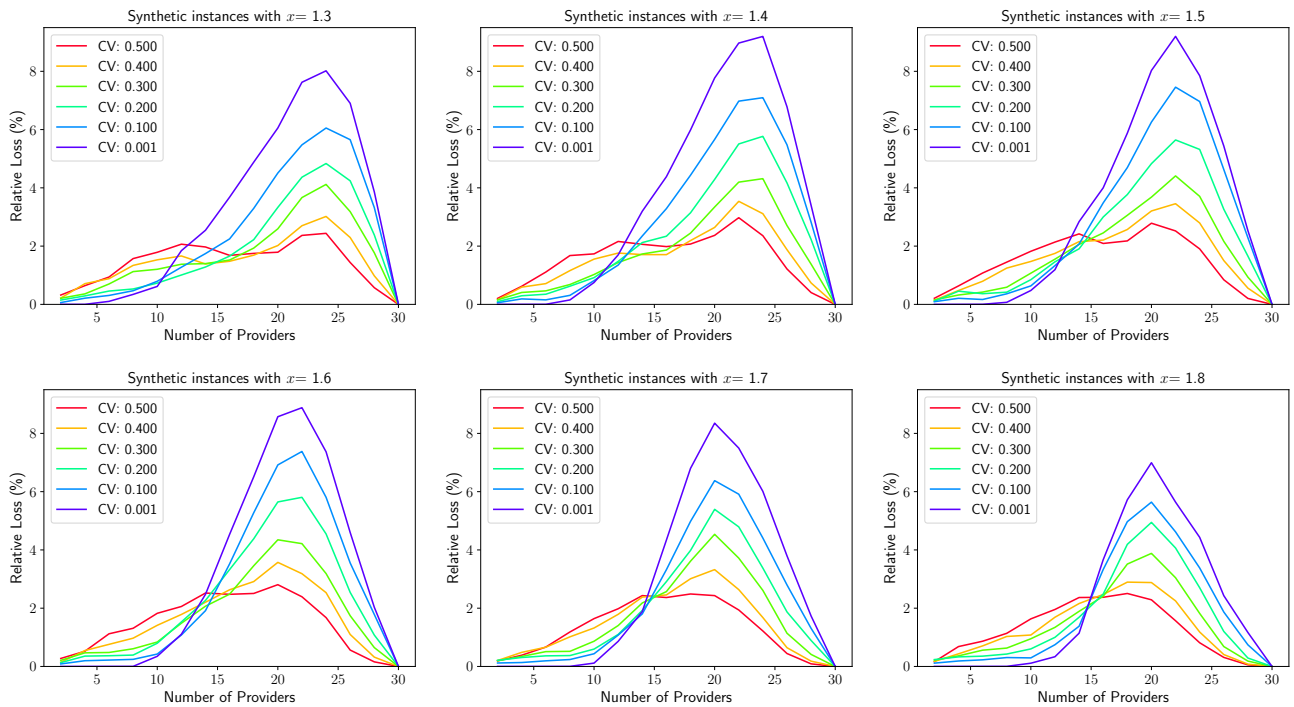

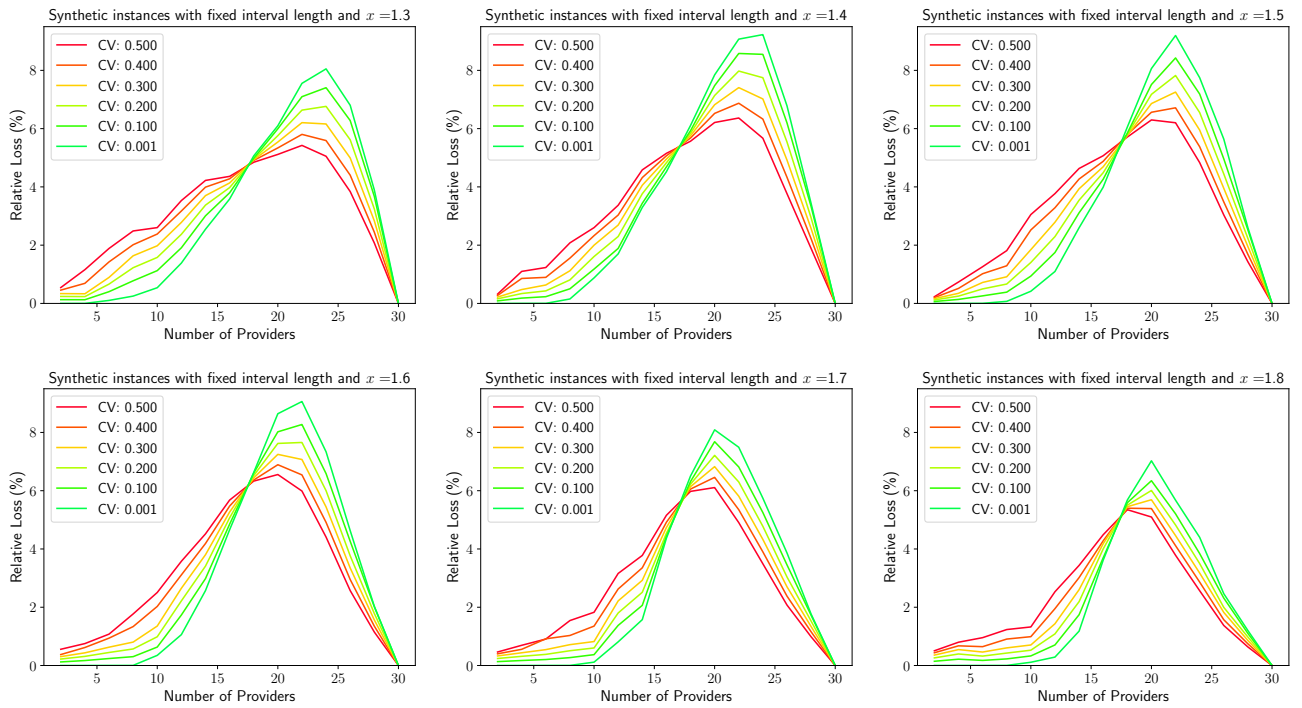Figure 13: **Synthetic instances for x varying from 1.3 to 1.8.**

Figure 14: **Synthetic instances with fixed interval length for x varying from 1.3 to 1.8.**