SOCIAL GOOD, FAIRNESS, AND EFFICIENCY IN OPERATIONS MANAGEMENT

A DISSERTATION
SUBMITTED TO THE GRADUATE SCHOOL OF BUSINESS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Xavier Sebastián Warnes

June 2021

This dissertation is online at: http://purl.stanford.edu/zb237jp6721

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Dan Iancu, Primary Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Yonatan Gur, Co-Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Erica Plambeck

Approved for the Stanford University Committee on Graduate Studies.

Stacey F. Bent, Vice Provost for Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format. An original signed hard copy of the signature page is on file in University Archives.

Abstract

In recent years, there has been an increasing interest in the Operations community in studying problems that balance the maximization of profits and efficiency with notions of fairness (e.g., Bertsimas et al. 2011, 2012) and environmental sustainability (e.g., Kleindorfer et al. 2005, Lee and Tang 2018). In this thesis, we present two works that contribute to this growing literature. The first chapter, coauthored with Yonatan Gur and Dan Iancu, studies the trade-offs between efficiency and guarantees to providers that may arise from equity or fairness considerations. In the second chapter, coauthored with Dan Iancu and Erica Plambeck, we investigate how increasing smallholder farmers' welfare through intensification can affect tropical forest conservation. We provide a more detailed description of each chapter below.

Value Loss in Allocation Systems with Provider Guarantees. Many operational settings share the following three features: (i) a centralized planning system allocates tasks to workers or service providers, (ii) the providers generate value by completing the tasks, and (iii) the completion of tasks influences the providers' welfare. In such cases, the planning system's allocations often entail trade-offs between the service providers' welfare and the total value that is generated (or that accrues to the system itself), and concern arises that allocations that are good under one metric may perform poorly under the other. In this chapter we propose a broad framework for quantifying the magnitude of value losses when allocations are restricted to satisfy certain desirable guarantees to the service providers. We consider a general class of guarantees that includes many considerations of practical interest arising, e.g., in the design of sustainable two-sided markets, in workforce welfare and compensation, or in sourcing and payments in supply chains, among other application

domains. We derive tight bounds on the relative value loss, and show that this loss is limited for any restriction included in our general class. Our analysis shows that when many providers are present, the largest losses are driven by fairness considerations, whereas when few providers are present, they are driven by the heterogeneity in the providers' effectiveness to generate value; when providers are perfectly homogenous, the losses never exceed 50%. We study additional loss drivers and find that less variability in the value of jobs and a more balanced supply-demand ratio may lead to larger losses. Lastly, we demonstrate numerically using both real-world and synthetic data that the loss can be small in several cases of practical interest.

Improving Smallholder Welfare While Preserving Natural Forest: Intensification vs Deforestation. Increasing the welfare of smallholder farmers in developing countries plays a crucial role in the global effort to reduce worldwide poverty and hunger. On the one hand, smallholders represent a large proportion of the world's poor and, on the other, they produce the majority of the food consumed in developing countries. This realization has led governments and organizations around the world to implement policies aimed at increasing farmers' yields. Although most of these policies have resulted in welfare increases, the environmental effects have been varied. While in many settings intensification policies have been linked to a decrease in deforestation, in many other settings the reverse is true. In this chapter we propose a novel explanation of these seemingly contradictory results. We achieve this through studying a detailed operational model of a farmer's dynamic decisions of land-clearing and production. We show the importance of considering the interaction between random production costs and liquidity constraints faced by smallholder farmers. These two elements are key to our main result: a reduction in the cost of intensification can lead to *lower* deforestation rates when the variation in production costs is high enough compared to the cost of intensification. Alternatively, the same reduction in the cost of intensification may lead to *higher* deforestation rates if the variation in production costs is low enough compared to the cost of intensification. This result helps explain the discrepancies seen in practices and may allow policy makers to better target interventions in order to achieve win-win situations: improvement of smallholder welfare and protection of the natural forest.

Acknowledgments

I am incredibly grateful to many people that have helped me over the years. I am grateful to Dan Iancu, for his support and advice, for agreeing to meet as many times as I needed, for always being there to help me grow as a researcher, for helping me find the questions that drive me to keep on working every day, for believing in me. I am grateful to Yonatan Gur, for pushing me to succeed, for his sharp and clear comments that constantly improved my research process and my writing, for his lessons on how to engage in research, for always responding to my emails immediately, even in the middle of the night. I am grateful to Erica Plambeck, for always sharing her deep insights and perspectives, for being the most positive and heartwarming person I've ever worked with, for her contagious drive to improve the world. Without their support and hard work this thesis would not have been possible.

To my fellow GSB students, who made coming to *the cubes* a blast, I am deeply grateful for our friendship. This last year has really made me value how much I depended on those day to day interactions with such an incredible group of people. To Nacho and Greg, I wouldn't have lasted a quarter without your support, help, and friendship. To Keka and Bree, for completing those other two, and for your friendship over these years. To Ilya, for his advice on how to be a better researcher, and for not letting me fall to my death (literally). To Chloe, Jonathan, and Fede, for all those movie nights, and great moments together. To Ashu and Jae, for all those long chats in *the cubes*, that made going to work such a pleasure. To Bar, for helping me prepare the *quals* and for never letting me win at tennis. To the OIT team: Danqi, Zhengli, Wenjia, Wanning, Mingxi, Ilan, Mine, and Bryce for making the OIT such a great group to be in.

I want to thank my parents, Gabriela and Ignacio, that made me who I am today. Without your example I would have never thought to start this journey, and without your unconditional support, I wouldn't have been able to succeed. I want to thank my brother, Pablo. Even through the physical distance separating us, your success and your drive will alway push me to succeed. I want to thank Marcelo and Silvia, for being the most supportive in-laws I could have asked for.

Most important of all, I want to thank my wife and companion in this journey, Sabri. I cannot express how much I relied on your love and support over all these years. I would not be complete without you.

Contents

Chapter 1

Value Loss in Allocation Systems with Provider Guarantees

1.1 Introduction

In many operational settings, a centralized planning system (*system* henceforth) decides how to allocate a pool of resources or tasks to an existing set of workers or service providers. Such allocations routinely determine the total value created; but oftentimes, they also influence how this value is shared between the system and the providers. When only some allocations meet particular attributes that the providers find appealing, this could significantly impact both the providers' welfare as well as the system in the long-run, e.g., 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 Upwork, Grubhub, Uber, or Lyft. These platforms match customer service requests for labor, food, or rides, with dedicated service providers (freelance workers or drivers), and typically rely on revenue-sharing agreements to split the revenue collected between the providers and the platform. The allocation of service requests thus critically influences the value generated upon each service completion as well as the portions of the total value that are retained by different providers. This value may consist of both monetary components (e.g., the revenue collected, or the profits when accounting for costs) as well as 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., GrubHub 2019, Uber Technologies Inc. 2018b and Lyft Inc. 2018) to designing loyalty and bonus payments tied to completing multiple service requests (see, e.g., DoorDash 2019, Financial Times 2014, and Uber Technologies Inc. 2018a). How much are platforms sacrificing in terms of revenues, profits, or customer experience when their allocations are designed to carry such guarantees for the service providers?

To demonstrate the 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 may also increase 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 terms, 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.

1.1.1 Main Contribution

On the modeling front, we develop a broad framework that allows quantifying the 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 (or resources) to a set of providers who are endowed with heterogeneous effectiveness to generate value by completing the jobs. We capture an allocation design that institutes desirable provider guarantees by imposing constraints (also referred to as *restrictions*) on the system's feasible allocations, and we only require the set of constrained allocations to satisfy a mild and natural monotonicity condition. This allows us to capture a variety of practical considerations arising in the sustainable design of two-sided markets, in workforce welfare and compensation mechanisms, as well as in sourcing and payments in supply chains, among other application domains.

We define the relative value loss associated with instituting certain provider guarantees as the fraction of the maximal value achievable by unrestricted allocations that is lost when imposing the constraints associated with satisfying these guarantees. We derive tight bounds on the relative value loss that hold for *any* restrictions in a general class of provider

[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.

guarantees. We establish these by solving a fractional linear relaxation of the problem of maximizing the relative value loss, and by producing instances that match the maximal value of this relaxation.

Our bounds only depend on the heterogeneity in the providers' effectiveness to generate value, and on the number of providers. We find that although unbounded heterogeneity may lead to unbounded relative loss, when heterogeneity is bounded, the loss must be bounded as well under any of the guarantees that we consider. In particular, the relative loss never exceeds 50% when providers are homogenous in their effectiveness to generate value, irrespective of how many providers exist. This contrasts several findings in the resource allocation literature that exhibit unbounded losses, particularly in regimes with a large number of providers, and highlights the importance of explicitly capturing certain aspects of the job allocation problem that are included in our model (see §1.1.2 for a more detailed discussion). Qualitatively, our findings suggest that allocation systems concerned with large losses due to instituting provider guarantees may consider reducing the heterogeneity in the providers' (endowed) effectiveness to generate value, as this can substantially reduce the magnitude of such losses in the worst case.

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

In contrast, we show that when many providers with bounded heterogeneity exist, the primary loss driver is *lost demand*, i.e., unallocated jobs: the largest losses occur when the system is unable to allocate many of the jobs due to a combination of physical constraints (preventing jobs from being completed by the same provider) and fairness constraints related to the guarantees instituted for the providers. More precisely, we find that the worst-case losses emerge from Max-Min fairness considerations. Such fairness considerations are a strict subset of the guarantees that we consider; thus, our result provides additional context for a recent literature stream focused on quantifying fairness-efficiency trade-offs in operations (see §1.1.2 for a more detailed discussion), by showing that these may become particularly pertinent in large-scale allocation systems with ample supply, ample demand and physical constraints that prevent certain jobs from being assigned together.

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

Using both synthetic data and real-world data consisting of taxi trips in New York City, we study numerically the relative value losses associated with implementing provider guarantees. As this setting corresponds roughly to viewing providers (i.e., taxi drivers) as approximately homogeneous in their effectiveness to generate value from jobs, we focus on guarantees corresponding to fairness considerations. The numerical results confirm the robustness of our theoretical findings concerning the impact of the variation in the intrinsic

[3]Nevertheless, we show numerically that the effect of asymmetric sets of feasible allocations is nuanced: some asymmetry may actually reduce average losses, but complete asymmetry can lead to unbounded worst-case losses.

job value and the supply-demand imbalance on the relative value loss. Moreover, in the instances generated from the real-world data, we document relative value losses that never exceed 4%. This suggests that losses associated with implementing provider guarantees in particular cases of practical interest may actually be even significantly lower than our theoretical worst-case bounds – a finding that we believe motivates additional research focused on more specific settings or on finding policies that can achieve such guarantees.

1.1.2 Related Literature

Relative Efficiency Losses due to Fairness. Our work is related to a stream of literature that studies efficiency losses emerging when outcomes are constrained to satisfy certain fairness considerations. Bertsimas et al. (2011, 2012) consider continuous resource allocation problems where a centralized decision maker balances efficiency (i.e., social welfare) with fairness and equity considerations. They define the *price of fairness* as the relative loss in efficiency under such fairness considerations, and derive theoretical bounds on this measure that depend on the number of agents and on the fairness criterion. As the number of agents grows arbitrarily large, the worst-case losses always approach 100%, regardless of the fairness criterion imposed, and the instances that get close to this loss involve asymmetries between the different agents. Our study retains the same definition of relative value loss, but differs in several modeling primitives and results. Specifically, we focus on *discrete* allocation problems where guarantees (such as fairness considerations) are modeled through constraints on the feasible allocations rather than on the possible utility outcomes; we consider a broader set of restrictions that includes the fairness criteria in Bertsimas et al. (2011, 2012) as special cases; and we explicitly model and study a specific form of heterogeneity, in the providers' effectiveness to generate value from the jobs they are assigned. In our setting, unbounded heterogeneity is needed to obtain losses of 100%, and this can occur with both a small and a large number of agents/providers, unlike in Bertsimas et al. (2011, 2012). More importantly, we find that when heterogeneity is bounded, the worst-case losses are always bounded irrespective of the number of providers and of the guarantees used, and never exceed 50% when providers are homogenous. We elaborate more on the root causes for this discrepancy in §1.4. Our finding that worst-case restrictions actually correspond to Max-Min fairness considerations when the number of agents/providers is sufficiently large also provides more context for the large stream of literature focused on quantifying fairness and efficiency trade-offs in various large-scale operational settings (e.g., Bertsimas et al. 2013, McCoy and Lee 2014, Iancu and Trichakis 2014, Qi 2017).

Efficiency Losses in Equilibrium. Our paper is also related to a rich literature studying the *Price of Anarchy* – a measure introduced by Papadimitriou (2001) and Roughgarden and Tardos (2002) that quantifies the efficiency loss of Nash equilibrium outcomes relative to an optimal centralized solution. It is known that the Price of Anarchy can be bounded in certain settings (see, e.g., Roughgarden 2003, 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). We study efficiency losses generated by a centralized planner restricting the outcomes on purpose, rather than losses due to strategic behavior 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, whereas our paper is focused on quantifying the losses associated with such allocations, and understanding their key drivers.

Efficiency of Contracts. Our work is also related to a body of literature studying the efficiency losses that may arise in various principal-agent interactions, such as between a firm's shareholders, debtholders and managers (see, e.g., Jensen and Meckling 1976), between firms and their sales associates (see, e.g., Farley 1964), between buyers and their suppliers (see, e.g., Cachon and Lariviere 2005), etc. Several papers in this literature also seek to quantify the associated efficiency losses; see, e.g., Besbes et al. 2017 for more discussion and additional references. Our work is not concerned with agency considerations, but instead focuses on quantifying the relative efficiency loss associated with restricting allocations to satisfy certain attributes that may be desirable to providers.

1.2 Problem Formulation

For the sake of clarity, we first provide a basic description of our setup, and then discuss some concrete examples in §1.2.1. A discussion of the modeling assumptions is deferred to §1.2.2.

Consider a centralized planning system that allocates a given set of jobs D to a set of n service providers denoted by $N = \{1, \ldots, n\}$. 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 $\boldsymbol{A} = (A_1, \ldots, A_n) \in \mathcal{F}$ denotes a feasible allocation, then A_i denotes the jobs allocated to provider $i \in N$, and $A_{-i} \stackrel{\text{def}}{=} (A_1, \ldots, A_{i-1}, A_{i+1}, \ldots, A_n)$ 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 $0 < \gamma_{\min} \leq \gamma_{\max}$[4]. We let $\boldsymbol{\gamma} \in [\gamma_{\min}, \gamma_{\max}]^n$ denote the heterogeneity profile of providers, and we denote the degree of heterogeneity by

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

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

[4]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.

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 $\mathcal{F}_{\mathrm{G}}$:

$$\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) = \frac{\max_{\boldsymbol{A}\in\mathcal{F}} \sum_{i=1}^{n} \gamma_i v(A_i) - \max_{\boldsymbol{B}\in\mathcal{F}_{\mathrm{G}}} \sum_{i=1}^{n} \gamma_i v(B_i)}{\max_{\boldsymbol{A}\in\mathcal{F}} \sum_{i=1}^{n} \gamma_i v(A_i)}. \tag{1.1}$$

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.

1.2.1 Examples

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

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

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

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

(iii) Costs/profits. The spatial and temporal length of a job are typically key drivers for both the revenues as well as the costs from completing the job. Thus, $\gamma_i v(d)$ could capture the gross profit when driver i completes the job, equal to the revenue $v(d)$ net of provision costs $(1 - \gamma_i)v(d)$. This allows modeling heterogeneity in the transportation costs, for example, 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 jobs, a measure of economic efficiency in the presence of costs.

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

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

Workforce Management. A closely related example occurs when managing a workforce such as a team of sales associates. In this case, the system could be a regional or store manager who designs schedules for n sales associates. These schedules could be temporal (e.g., which hours or days to work), spatial (e.g., which floors or departments to cover) or could consist of the assignments of particular clients. With $v(d)$ denoting the (expected) revenue from a particular sales opportunity d, $\gamma_i v(d)$ can capture the revenue generated when this opportunity is assigned to associate i (with γ_i measuring the associate's effectiveness/performance), or it could capture the fraction of the revenue accruing to the firm (with $1 - \gamma_i$ denoting associate i's commission). As in the service platform example, the total value could also capture the quality of service experienced by clients or the gross profit when the associates incur variable costs; the set of feasible allocations $\mathcal{F}$ could capture constraints on the schedules or assignments; and the subset of allocations $\mathcal{F}_{\mathrm{G}} \subseteq \mathcal{F}$ 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, effectiveness, or quality) or gross margins for the firm (when different prices are paid to different suppliers). The set of allocations $\mathcal{F}_{\mathrm{G}} \subseteq \mathcal{F}$ 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).

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

Table 1.1: **Interpretation of parameters.** For each example in §1.2.1 we present a summary of the interpretation of each parameter.

1.2.2 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 1.1 *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* $i \in N$.

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

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

Part (iii) states that it is always possible to allocate *fewer* jobs. This always holds when jobs can be carried out independently from one another. In addition, the requirement allows certain dependencies between jobs: for instance, if a set of jobs must be completed together, i.e., by a single provider, these could be grouped into a single aggregate job that should be allocated as an indivisible object in our setup.

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

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

Assumption 1.2 (Pareto-efficient Provider Guarantees) *If* $\boldsymbol{B} \in \mathcal{F}_{\mathrm{G}}$ *and* $\boldsymbol{A} \in \mathcal{F}$ *are such that* $v(A_i) \geq v(B_i)$ *for all* $i \in N$*, then* $\boldsymbol{A} \in \mathcal{F}_{\mathrm{G}}$*.*

This assumption bears an immediate interpretation as a requirement of Pareto efficiency: it asks that any provider guarantee (as captured by $\mathcal{F}_{\mathrm{G}}$) should allow for more value to be generated whenever that is possible; if an allocation satisfies the provider guarantees, then feasible allocations where every provider can generate *more* value should also satisfy these guarantees. Seen in this light, the assumption comes across as a natural axiomatic requirement that should be satisfied by any valid guarantees that the system might offer its providers.

In a different sense, if the set $\mathcal{F}_{\mathrm{G}}$ captures allocations that are "acceptable" to the providers (e.g., as they carry desirable guarantees), then Assumption 1.2 can also be interpreted as a condition on the providers' preferences: namely, that assigning higher value jobs to all providers should not be "less acceptable." This suggests that the provider preferences encoded in the set of guarantees should be aligned with the value function v. The following result, which is an equivalent characterization of Assumption 1.2, formalizes this intuition.

Proposition 1.1 **($\mathcal{F}_{\mathrm{G}}$ Aligned with v)** *A subset $\mathcal{F}_{\mathrm{G}} \subseteq \mathcal{F}$ satisfies Assumption 1.2 if and only if $\mathcal{F}_{\mathrm{G}}$ satisfies:*

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

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.1. The result establishes that all valid provider guarantees correspond to value-maximizing allocations when jobs are evaluated according to some function g that preserves the same ordering as the value function v. Insofar as allocations with guarantees $\mathcal{F}_{\mathrm{G}}$ capture provider preferences, this also means that Assumption 1.2 exactly requires such preferences to be "aligned" with v. As we demonstrate in the next subsection, such alignments between captured value and provider guarantees (and preferences) can arise naturally in many practical settings.

1.2.3 Discussion and Classes of Provider Guarantees

Income Guarantees under Monotonic Payment Functions

An important class of provider guarantees satisfying Assumption 1.2 arises from ensuring a minimum level of total income to each provider, when every provider's compensation is increasing in the value captured by job completions. More precisely, suppose that a provider who completes a set of jobs $S \subseteq D$ is compensated with an amount $p(S)$, where $p : \mathcal{P}(D) \to \mathbb{R}$ denotes a payment function. For a given real number τ, we define the set

$$\mathcal{F}_{\mathrm{G}}(p, \tau, N) := \{\boldsymbol{A} \in \mathcal{F} \mid p(A_i) \geq \tau, \ \forall\, i \in N\}, \tag{1.2}$$

that includes the allocations that ensure that each provider is compensated with at least an amount τ. The set $\mathcal{F}_{\mathrm{G}}(p, \tau, N)$ satisfies Assumption 1.2 if p satisfies the property:

$$v(S) \geq v(T) \quad \Rightarrow \quad p(S) \geq p(T), \ \forall\, S, T \in \mathcal{P}(D). \tag{1.3}$$

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

The class of minimal income guarantees could also be generalized by considering provider-specific payments and guarantees, i.e., by taking p_i or τ_i or by considering guarantees only for a subset $\hat{N} \subseteq N$ of the providers. A visual depiction of several such guarantees is shown in Figure 1.1, for the revenue-sharing case with $p(S) = \beta v(S)$ for $\beta = 1$. The figure shows both *uniform income guarantees* (with a unique τ for all providers), as well as 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 $\mathcal{F}_{\mathrm{G}}$ 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 §1.2.3.

Unions and Intersections of Guarantees

Consider any collection of sets $\{\mathcal{F}_G^k\}_{k\in K}$ where each set $\mathcal{F}_G^k$ satisfies Assumption 1.2. Then, $\cap_{k\in K}\mathcal{F}_G^k$ and $\cup_{k\in K}\mathcal{F}_G^k$ also satisfy Assumption 1.2. Considering intersections is useful for modeling restrictions that certain targeted providers find desirable or acceptable. Namely, building on the intuition developed in Proposition 1.1, suppose each provider $i \in N$ is

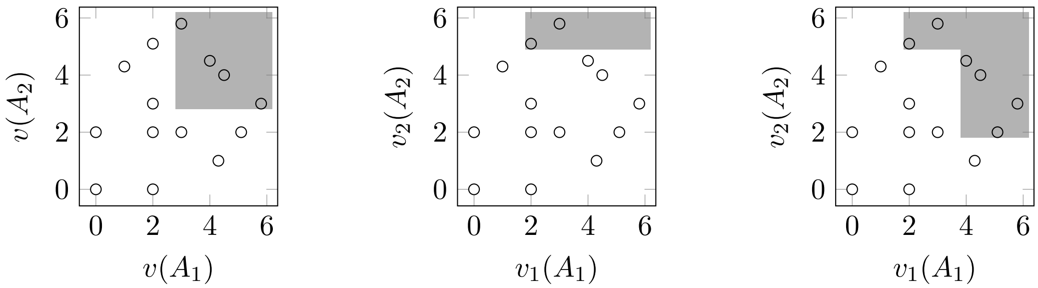

Figure 1.1: **Examples of provider guarantees.** Income guarantees $\mathcal{F}_{\mathrm{G}}(p, \tau, N)$ for two providers ($n = 2$) compensated according to revenue-sharing agreements with $p(S) = \beta v(S)$ for $\beta = 1$. Each circle denotes a feasible allocation, with the two axes corresponding to the intrinsic values $v(A_1)$ and $v(A_2)$ for each provider. The circles in the shaded area show the revenues achievable by allocations in each $\mathcal{F}_{\mathrm{G}}$. **(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.

endowed with a utility function u^i satisfying the co-monotonicity requirements in Proposition 1.1.[5] In 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 $\mathcal{F}_{\mathrm{G}} = \cap_{i \in \hat{N}} \mathcal{F}_G^i$ as the allocations with guarantees. Alternatively, we could also consider a "satisficing" model (see Simon 1956) where $\mathcal{F}_G^i := \{\boldsymbol{A} \in \mathcal{F} \,:\, u^i(\boldsymbol{A}) \geq \tau_i\}$ 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 1.1. In fact, the following equivalent characterization of Assumption 1.2 shows that *any* provider guarantee satisfying it can actually be written as the union of income guarantees under monotonic payment functions.

Proposition 1.2 *The set of allocations with guarantees $\mathcal{F}_{\mathrm{G}}$ satisfies Assumption 1.2 if and*

[5] In particular, $u(\boldsymbol{B}) \geq u(\boldsymbol{A})$ for any $\boldsymbol{A}, \boldsymbol{B} \in \mathcal{F}$ with $v(B_i) \geq v(A_i)$, $\forall\, i \in N$.

only if it can be expressed as the union of income guarantees under monotonic payment functions. That is, $\mathcal{F}_{\mathrm{G}} \subseteq \mathcal{F}$ *satisfies Assumption 1.2 if and only if there exist monotonic payment functions and income guarantees* $\{(\boldsymbol{p}_k, \boldsymbol{\tau}_k)\}_{k\in K}$ *for some index set* K *such that* $\mathcal{F}_{\mathrm{G}} = \cup_{k\in K} \mathcal{F}_{\mathrm{G}}(\boldsymbol{p}_k, \boldsymbol{\tau}_k, N)$*.*

Proposition 1.2 implies that the class of income guarantees under monotonic payment functions is in some sense a universal generating family for all the restrictions that satisfy Assumption 1.2, as any such restriction can be captured by considering several alternatives from the former generating family. Although this characterization is not directly used in the proof of the following results, it adds to the understanding of the class of guarantees that satisfy Assumption 1.2. A proof of Proposition 1.2 can be found in Appendix A.1.

Fairness

Assumption 1.2 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:

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

Because this welfare function is increasing in each component A_i, the associated restriction $\mathcal{F}_{\mathrm{G}} = \arg\max_{\boldsymbol{A}\in\mathcal{F}} W_\alpha(\boldsymbol{A})$ satisfies Assumption 1.2 (this is an immediate corollary of Proposition 1.1). The constant elasticity social welfare function is concave and component-wise increasing, and thus it exhibits diminishing marginal welfare increase as the values generated increase. In other words, if $v(A_i) < v(A_j)$, for providers i and j, then a marginal increase in $v(A_i)$ would lead to a larger welfare increase than a marginal increase in $v(A_j)$.

Additionally, the rate at which marginal increases diminish is governed by the parameter α; at higher α, there are higher incentives to increase the value allocated to providers that are worse off, which makes solutions "more fair." This motivates the name of α as the *inequality aversion parameter*; in fact, maximizing $W_\alpha(\boldsymbol{A})$ for $\alpha = 0$ recovers an efficient solution, while $\alpha \to \infty$ recovers *Max-Min fair* solutions (see Kalai and Smorodinsky 1975, Mas-Colell et al. 1995).

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

$$\mathcal{F}_{\mathrm{G}}^{\mathrm{mM}}(p, N) := \mathcal{F}_{\mathrm{G}}(p, \tau_{\max}, N), \qquad \text{where} \qquad \tau_{\max} := \max\{\tau \mid \mathcal{F}_{\mathrm{G}}(p, \tau, N) \neq \emptyset\}. \tag{1.4}$$

1.3 Bounding the Value Loss under Provider Guarantees

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

Theorem 1.1 (Upper bound on the value loss) *The value loss under provider guarantees is bounded above as follows:*

$$\sup_{\mathcal{F}, \mathcal{F}_{\mathrm{G}}, \boldsymbol{\gamma}} \mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) \leq \max\left\{\delta, \frac{n-1}{n+(1-\delta)(n-1)}\right\}, \tag{1.5}$$

where the supremum is taken with respect to all $\boldsymbol{\gamma} \in [\gamma_{\min}, \gamma_{\max}]^n$ and all sets $\mathcal{F}$, $\mathcal{F}_{\mathrm{G}}$ satisfying Assumptions 1.1 and 1.2.

Theorem 1.1 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 1.2 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.

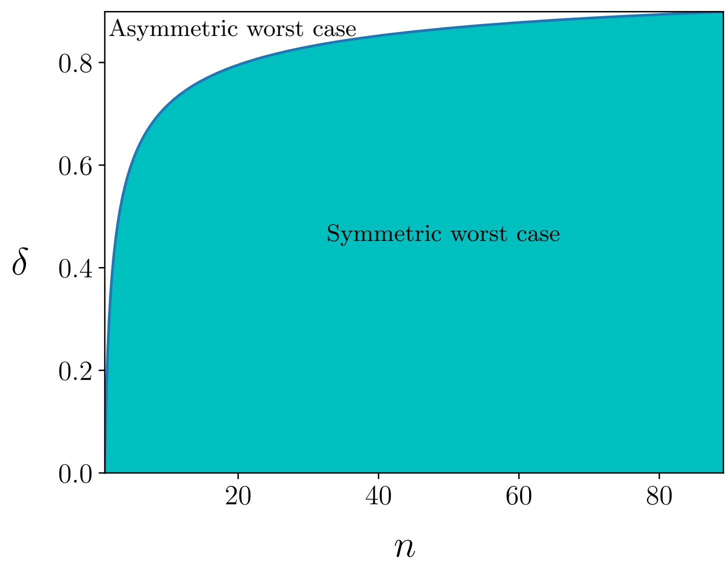

Figure 1.2: **Different types of worst cases in different parametric regions.** Shaded area denotes values of (n, δ) such that $\delta \leq \frac{n-1}{n+(1-\delta)(n-1)}$. In the non-shaded area, the worst-case instance corresponds to asymmetric guarantees (see Instance 1.1), where all jobs are guaranteed to only one provider, while in the shaded area the worst-case instance corresponds to a symmetric guarantee (see Instances 1.2 and 1.3), where $\mathcal{F}_{\mathrm{G}}$ corresponds to a Max-Min fairness guarantee.

The dependency of the bound on n and δ is depicted in Figure 1.3. It can be seen from (1.5) 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 §1.1.2, this result distinguishes our findings from many studies that document unbounded losses in contexts of price of fairness (Bertsimas et al. 2011) or price

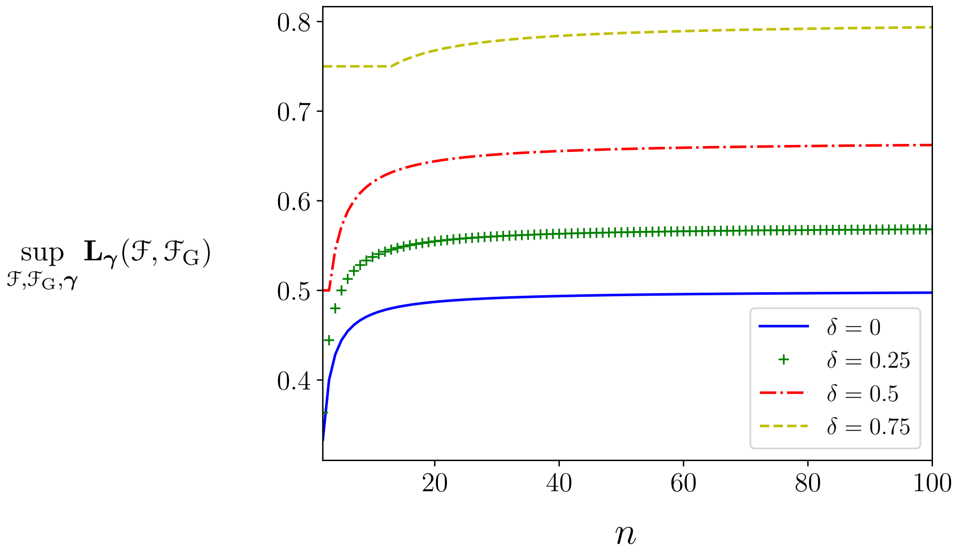

Figure 1.3: **Worst-case loss as a function of the number of providers and the heterogeneity level.** Plot of $\max\left\{\delta, \frac{n-1}{n+(1-\delta)(n-1)}\right\}$ as a function of n for different values of δ.

of anarchy (Koutsoupias and Papadimitriou 2009).

Main ideas in the proof. We defer the complete proof of Theorem 1.1 to Appendix A.1, but we briefly describe its main ideas here. First, we propose an LP relaxation of the problem of maximizing $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}})$ over all sets $\mathcal{F}$ and $\mathcal{F}_{\mathrm{G}}$, for a fixed vector $\boldsymbol{\gamma}$. For this purpose, we provide a family of inequalities that connect the maximum value over the set of feasible allocations $\mathcal{F}$ and over the set of allocations with guarantees $\mathcal{F}_{\mathrm{G}}$. 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 1.1 is in fact tight, by characterizing instances and provider guarantees that achieve the worst-case loss.

Theorem 1.2 (Attainable value loss) *The bound on value loss in Theorem 1.1 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 1.1-1.2 so that

$$\mathbf{L}_{\gamma^1}(\mathcal{F}^1, \mathcal{F}^1_G) = \delta \tag{1.6}$$

$$\mathbf{L}_{\gamma^2}(\mathcal{F}^2, \mathcal{F}^2_G) = \frac{n-1}{n + (1-\delta)(n-1)} - \varepsilon. \tag{1.7}$$

Theorem 1.2 states that the worst-case relative losses characterized in Theorem 1.1 are tight. To prove this result, we exhibit three classes of instances, where the first one achieves the loss in (1.6) and the other two asymptotically achieve the loss in (1.7). To provide more intuition, we describe these instances for the case with $n = 2$ here, and defer the general case to Appendix A.1.

Instance 1.1 (High guarantees with high provider heterogeneity) *Consider $n = 2$ providers with $\gamma_1 = \gamma_{\max}$ and $\gamma_2 = \gamma_{\min}$, a set of jobs D with $v(D) > 0$, a set of allocations with guarantees $\mathcal{F}_{\mathrm{G}} = \{(\emptyset,\ D)\}$, and any feasible set $\mathcal{F}$ with $\mathcal{F}_{\mathrm{G}} \subset \mathcal{F}$. Then, $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) = \frac{\gamma_{\max} - \gamma_{\min}}{\gamma_{\max}} = \delta$.*

In Instance 1.1, the system has a set of jobs D to assign to two providers, and it is possible to assign all the jobs to a single provider.[6] The first provider generates more value for the system ($\gamma_1 \geq \gamma_2$), so the value-maximizing allocation would be $(D,\ \emptyset)$, assigning all jobs to that provider. However, the only allocation with guarantees is $(\emptyset,\ D)$, requiring all jobs to be assigned to the second provider, and thus generating a loss due to heterogeneity. This results in worst-case losses when heterogeneity is high and the number of providers is not too large, as depicted in Figure 1.2.

The high losses in Instance 1.1 are enabled by the joint presence of two key features: a high degree of heterogeneity in the providers' effectiveness to generate value (with one provider maximally effective, and another minimally effective), and asymmetric guarantees that allocate all the jobs to the least effective provider. This asymmetry is critical: if the guarantees had been "symmetrized," e.g., by considering the union of all permutations of

[6] Note that by symmetry, since $(\emptyset,\ D) \in \mathcal{F}$, we must also have $(D,\ \emptyset) \in \mathcal{F}$, so this allocation is also feasible.

which provider receives the full set of jobs, then the loss would vanish. So for heterogeneity to be a critical driving force, it must be accompanied by asymmetric guarantees that are misaligned in order to force jobs being assigned to less effective providers. The inefficiency in Instance 1.1 can thus play a particularly prominent role in service platforms that may favor new providers by guaranteeing them a higher workload, when such providers also have less experience and thus generate less value, for example, in terms of productivity as well as goodwill gain and customer satisfaction.[7]

The following instance achieves the bound on the relative loss that is given in (1.7).

Instance 1.2 (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 1.4: 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 $\mathcal{F}_{\mathrm{G}}^{\mathrm{mM}}(p, N)$ under any strict monotonic payment function p, as described in (1.4). Then, the value-maximizing allocation is $\big(\{d_1, d_2\}, \{d_3\}\big)$, the value-maximizing allocation with guarantees is $\big(\{d_1\}, \{d_2\}\big)$,[8] 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

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

Instance 1.2 describes a situation where the system allocates jobs that have a certain

[7] 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).

[8] 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 $v(d_3) = 1 - \kappa < 1 = v(d_1)$.

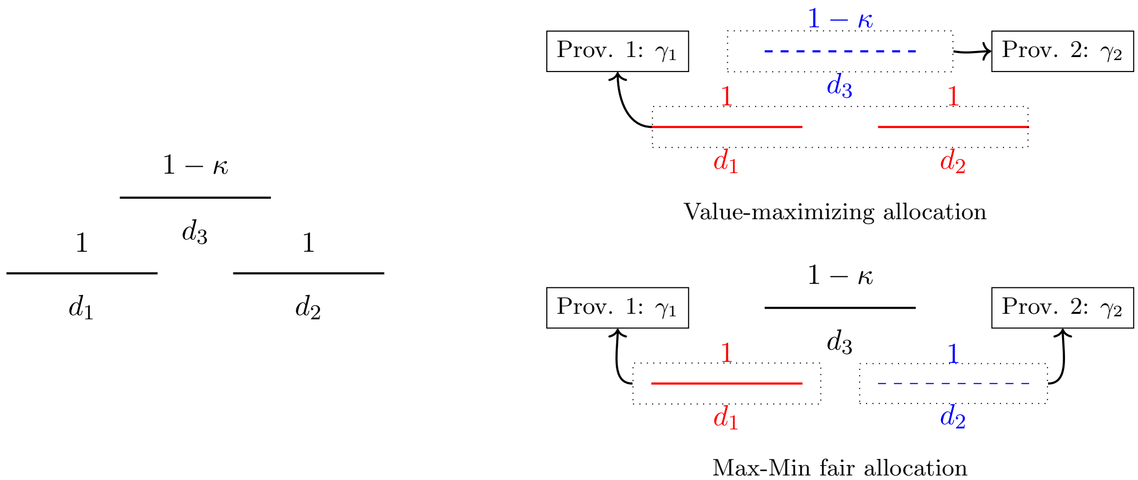

Figure 1.4: **Symmetric worst-case instance.** The left panel depicts Instance 1.2 that consists of $n = 2$ providers and three jobs $D = \{d_1, d_2, d_3\}$ represented by segments. Jobs can be assigned together if the corresponding segments are non-overlapping. The right panel depicts the value-maximizing allocation at the top, and the Max-Min fair allocation at the bottom.

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 1.2 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, d_2.

Instance 1.2 thus showcases two new drivers for the value loss. First is the structure of the set of feasible allocations $\mathcal{F}$, which contains certain exclusion constraints that force job d_3 to be incompatible with both jobs d_1 and d_2. Note that if job d_3 could be allocated together with either of the two other jobs, then the loss $\mathbf{L}_\gamma(\mathcal{F}, \mathcal{F}_G)$ would vanish. We revisit such exclusion constraints in §1.4, where we analyze their impact on the value loss in more detail. Second is the variation in the intrinsic value of jobs; although the loss grows as jobs become more similar in value (i.e., as $\kappa \to 0$), some variation is in fact critical: if $\kappa = 0$, the loss $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}})$ would again vanish. Our next instance shows that this variation is actually *not* critical to achieving a worst-case loss.

Instance 1.3 (Max-Min fairness with monotonic payments and equal-valued jobs) *Consider* $n = 2$ *providers with* $\gamma_1 = \gamma_{\max}$ *and* $\gamma_2 = \gamma_{\min}$*, and a set of jobs* $D = \{d_1, d_2, d_3, d_4, d_5\}$ *with* $v(d_i) = 1$ *for each* $i \in N$*. The feasible allocations are depicted in Figure 1.5: 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 (1.4). 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* $\big(\{d_1, d_2\}, \{d_3, d_4\}\big)$*. Therefore,* $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) = \frac{4\gamma_{\max}+\gamma_{\min}-2(\gamma_{\max}+\gamma_{\min})}{(4\gamma_{\max}+\gamma_{\min})}$.[9]

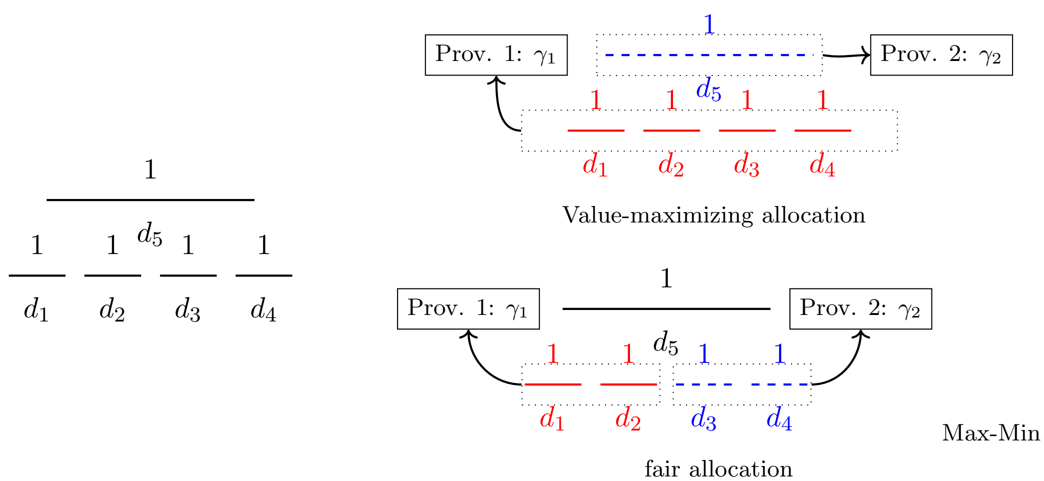

Figure 1.5: **Symmetric worst-case instance with equal-valued jobs.** The left panel depicts Instance 1.3: there are $n = 2$ providers and five jobs $D = \{d_1, d_2, d_3, d_4, d_5\}$ represented by segments. Jobs can be assigned together if the corresponding segments are non-overlapping. The right panel depicts the value-maximizing allocation at the top, and the value-maximizing Max-Min fair allocation at the bottom.

Instance 1.3 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

[9] As in Instance 1.2, 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.

construction are provided in Instance A.3 of Appendix A.1). This generalization yields a worst-case loss of:

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

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

Instance 1.3 shares certain similarities with Instance 1.2. Both rely on having a large number of providers and jobs, and on the presence of exclusion constraints in the set of feasible allocations $\mathcal{F}$ that prevent certain jobs from being allocated together. Also, the worst-case guarantees in both instances correspond to symmetric Max-Min fairness guarantees obtained under any strictly monotonic payment function. This suggests that in allocation systems with many providers and sufficiently "standardized" jobs (so that provider heterogeneity is not too large), symmetric fairness guarantees can be critical drivers of loss, as depicted in Figure 1.2. Lastly, it is worth emphasizing that although both Instance 1.2 and Instance 1.3 involve some heterogeneity in the providers' effectiveness, this is not critical to identifying these as worst-case instances; in fact, both instances yield worst-case losses when providers are homogeneous ($\delta = 0$).

Instance 1.3 also exhibits some notable differences from Instance 1.2: it relies on jobs with identical value, and it requires an arbitrarily large number of jobs to achieve the worst-case loss. We return to discuss each of these key drivers — the variation in job values and the supply/demand imbalance — and their relationship in more detail in §1.4. Additionally, note that the value of jobs in Instance 1.3 *cannot* be proportional to the size of the associated intervals; in contrast, Instance 1.2 can be constructed to satisfy this, which brings it closer to several practical applications, such as when intervals represent the time to complete a job and the value generated is proportional to this time.

We conclude our discussion by noting that all the loss drivers that are critical in Instances 1.2-1.3 are *irrelevant* in Instance 1.1, where the loss is entirely driven by the heterogeneity in provider effectiveness. To see this, note that all jobs in Instance 1.1 are assigned under all the allocations with guarantees $\mathcal{F}_{\mathrm{G}}$; thus, no value is lost due to jobs that remain unallocated as a result of physical constraints or restrictions imposed by the provider guarantees. Additionally, note that as long as jobs have some positive intrinsic value in Instance 1.1, any distribution of values could lead to worst-case losses; thus, the precise job values are completely irrelevant. Lastly, note that Instance 1.1 can be readily generalized to obtain the same loss in a setting with an arbitrary number of providers, as long as at least one maximally effective and one minimally effective provider exist and the guarantees require allocating jobs only to the latter (see Appendix A.1); thus, having a large number of providers or imbalanced supply/demand does not impact the loss in Instance 1.1.

1.4 Analysis of Key Loss Drivers

The previous section highlighted several potential drivers for value loss. Perhaps the most prominent of these is provider heterogeneity: when providers differ in their effectiveness to generate value (i.e., δ is large), this heterogeneity becomes the dominant loss driver, as evidenced in Instance 1.1. In this section we focus our discussion on several additional loss drivers: (*i*) the structure of the set of feasible allocations $\mathcal{F}$; (*ii*) the variation in the intrinsic values of jobs; (*iii*) the balance between supply (number of providers) and demand (number of jobs); and (*iv*) the integrality of allocations.

1.4.1 Structural Properties of the Set of Feasible Allocations

When heterogeneity is not the main loss driver, the structure of the set of feasible allocations $\mathcal{F}$ can be crucial, as Instances 1.2 and 1.3 demonstrated; recall that each of those instances involved certain exclusion constraints, whereby some jobs could not be allocated together to the same provider. Our next example shows that such constraints are critical: when the set $\mathcal{F}$ only includes constraints on how many jobs can be allocated together but without other explicit exclusion constraints, the value loss vanishes. Notice how we assume that all providers are homogeneous (all γ_i are equal) in order to eliminate the effect of heterogeneity on the value loss.

Instance 1.4 *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 $\mathcal{F} = \{(A_1, \ldots, A_n) \mid A_i \subseteq D, |A_i| \leq k_i, \forall i,$ and $A_i \cap A_j = \emptyset, \forall i \neq j\}$.*

Proposition 1.3 *Instance 1.4 satisfies $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) = 0$ for any set of allocations with guarantees $\mathcal{F}_{\mathrm{G}}$.*

This result becomes even more striking when the structure of the feasible sets in Instances 1.2 and 1.3 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 1.5 *Consider any $\boldsymbol{\gamma}$, a set of jobs $D = S \cup C$, and a set of feasible allocations of the form:*

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

The jobs described in Instance 1.5 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 1.4 (with capacity $k_i = \infty$ or $k_i = 1$, respectively), and thus result in zero loss. It is thus striking that by simply combining these two sets as done in Instance 1.5, one in fact obtains an instance that could either have *zero* loss or *a worst-case* loss, as formalized in the next result.

Proposition 1.4 *Consider Instance 1.5, then:*

(i). If $v(d) = v$, for all $d \in D$, and $|S| < 2n$, $\mathbf{L}_{\gamma}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) = 0$ for any uniform income guarantee $\mathcal{F}_{\mathrm{G}}$.

(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 1.2 by taking $\mathcal{F}_{\mathrm{G}}$ as the Max-Min fair allocations under any strictly monotonic payment function.

Proposition 1.4 implies that the feasible set structure can carry significant impact, but also that this structure in isolation is not a good predictor of value loss: the same structure could generate very large or very small losses, depending on other problem features such as the value of jobs or the imbalance between supply (number of providers) and demand (number of jobs). Notice how all conditions in (ii) satisfy the conditions in (i), except for the deviation in the intrinsic values of jobs in C. This also shows how small changes in an instance can lead to large changes in the value loss.

We conclude by noting that the symmetry of the set of feasible allocations $\mathcal{F}$ is also very important for our results. If this assumption were relaxed, e.g., if providers had different sets of jobs they could complete, then we can actually achieve a worst-case loss that asymptotically approaches 100% as the number of agents grows large (see Instance A.4 of Appendix A.1). This instance is actually inspired by the bandwidth allocation problem considered in Bertsimas et al. (2011), and matches the upper-bound on the price of fairness proved therein. This finding is consistent with our earlier result concerning the impact of heterogeneity in the providers' effectiveness to generate value from jobs, and it reinforces the idea that when providers are sufficiently "different," this can critically drive the losses, making them unbounded in extreme cases. To further explore this, in Appendix A.2 we analyze numerically the average relative loss for a family of instances with a varying degree of *symmetry*, measured by the proportion of all the jobs that each provider can perform. Interestingly, we find that as the instances become more asymmetric, the average losses decrease.

Intuitively, when each provider can perform a smaller fraction of all jobs, the probability of two providers being able to perform the same job decreases. In the extreme case this leads to zero loss: if no job can be performed by two providers, then the allocation problems can be separated into disjoint problems for each provider, which, by Proposition 1.5 implies that the value loss will be zero. Nevertheless, this decrease in the average loss is tied to the assumption that the subset of jobs a provider can perform is independent across providers. In fact, in Instance A.4 of Appendix A.1, the high loss is driven by a specific segmentation of the providers into two groups, each with a particular set of jobs they can perform. This suggests that the relative value loss is significantly influenced by the specific notion of symmetry considered, and that the impact of symmetry can be nontrivial: allowing for very high asymmetry might lead to very high worst-case losses, but asymmetry might also decrease the average losses in some cases.

1.4.2 Variation in Intrinsic Job Values

Instance 1.2 showed that the difference in the intrinsic value of jobs can be a key driver of the value loss, and that this loss can decrease as jobs have increasingly different intrinsic values. To further explore the impact of this feature, we now consider a slight modification of Instance 1.2 where we introduce a random perturbation in the value of one of the jobs, and we consider the expected loss as a function of the variance of this perturbation. Since each realization of the random perturbation corresponds to a particular Instance 1.2, when the variance of the perturbation is high we will obtain higher differences in the intrinsic values of jobs.

Example 1.1 *Consider Instance 1.2 with* $\gamma_{\min} = \gamma_{\max}$*, and assume that* κ *is uniformly distributed* $\kappa \sim U[-\frac{\Delta}{2}, \frac{\Delta}{2}]$*, so that* $v(d_3) \sim U[1 - \frac{\Delta}{2}, 1 + \frac{\Delta}{2}]$*. By taking the expectation of* *the loss with respect to this random variable, we get:*

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

Note that $\mathbb{E}[\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F},\mathcal{F}_{\mathrm{G}})]$ *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* $\mathbb{E}[\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F},\mathcal{F}_{\mathrm{G}})]$ *is decreasing in the variance of* $v(d_3)$*: higher variance implies lower expected loss.*

Interestingly, this example suggests that when intrinsic values are randomly generated, higher variance in values could actually *reduce* the impact of implementing provider guarantees, and result in a *lower* expected loss. The example can be generalized to a case with n providers (see Example A.1 in Appendix A.1), and we also confirm its robustness in a more realistic setting as part of our numerical exercise in §1.5. 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 1.2, 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 §1.5.

1.4.3 Supply-Demand Balance

Instance 1.3 suggested that the imbalance between supply (number of providers) and demand (number of jobs) can also critically drive the value loss. To build further intuition for this, we examine certain extreme cases that allow an analytical characterization. The next result shows that when a single provider or a sufficiently large number of providers are present, the value loss vanishes. We again limit the effect of heterogeneity on the value loss by assuming that all providers are homogeneous.

Proposition 1.5 *Assume that* $\gamma_{\min} = \gamma_{\max}$*, and consider any set of allocations with guarantees* $\mathcal{F}_{\mathrm{G}}$*. Then* $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) = 0$ *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 1.6 *Assume that* $\gamma_{\min} = \gamma_{\max}$*, and let* $n > |D|/2$ *and* $v(d) = v$ *for all* $d \in D$*. Then,* $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) = 0$*, for any uniform income guarantee* $\mathcal{F}_{\mathrm{G}}$*.*

The intuition behind this result is that when there are insufficient jobs to guarantee each provider at least two jobs, then there will be no value loss from guaranteeing the same level of income to all providers. This suggests that the value loss is likely small when the supply-demand imbalance is high, i.e., the ratio of supply to demand is either very low or very high. Although it is hard to analytically prove this more generally, we confirm it in our numerical tests in §1.5, 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.

1.4.4 Integrality of Allocations

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

$$\mathcal{F}^c = \left\{ \left(\{\theta_j\}_{j=1}^k, \{\boldsymbol{A}^j\}_{j=1}^k \right) \mid 0 \leq \theta_j \leq 1, \sum_{j=1}^k \theta_j = 1, \boldsymbol{A}^j \in \mathcal{F}, \forall j \in \{1, \ldots, k\}, k \geq 0 \right\}. \tag{1.8}$$

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

$$v(C_i) = \sum_{j=1}^{k} \theta_j v(A_i^j). \tag{1.9}$$

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

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

Proposition 1.7 *Given any* $\boldsymbol{\gamma}$ *and any set of feasible allocations* $\mathcal{F}$*, consider the set* $\mathcal{F}^c$ *defined in* (1.8) *and the extension of* $v(C)$ *defined in* (1.9)*. Then,* $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}^{\mathrm{c}}, \mathcal{F}_{\mathrm{G}}^{\mathrm{c}}) = 0$ *for any set of uniform income guarantees under monotonic payments* $\mathcal{F}_{\mathrm{G}}^{\mathrm{c}}$*.*

The intuition behind Proposition 1.7 is that when fractional allocations are possible, the system can simply consider an allocation obtained by mixing with an equal proportion $\frac{1}{n!}$ all the permutations of a particular value-maximizing allocation. This new allocation would still achieve the maximum value, while also allowing each agent to generate exactly the same value, and thus remaining feasible under any uniform income guarantee.

Our definition of fractional allocations is motivated by settings with finite periods and similar jobs being allocated in each period. Thus, while the jobs being allocated cannot be partitioned in each period, overall the allocations may be randomized across periods. Therefore, our result in Proposition 1.7 suggests that for certain provider guarantees defined with respect to average monotonic payment functions, the loss can be driven close to zero by correctly randomizing the allocations.

1.5 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 §1.2.3. To obtain the relative value loss $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}})$, 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 A.2.

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 ensures that part (*ii*) of Assumption 1.1 concerning the standardized nature of jobs is satisfied, in that any provider in the area can perform any subset of (non-overlapping) trips; in Appendix A.2, we present more complex feasibility restrictions that allow us to consider all trips in the city while still satisfying this Assumption. For each neighborhood, we focus on the first week of January 2016; we consider, for each particular day in that week, all the trips completed between 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 30 trips as subintervals of $(0, x]$, with the starting point of the trip sampled uniformly, and the trip length drawn from a truncated normal distribution; we use a mean of 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 A.2 for a more complete analysis. Figures 1.6 and 1.7 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 1.7, all values were taken to be equal. This structure resembles Instance 1.5, 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.

In all our instances, when all the job values are set equal – so that the variation in values disappears – we obtain a loss of zero. This is consistent with Proposition 1.4, and is reasonable to expect precisely because many of our data-driven instances match Instance 1.5, which is the premise for Proposition 1.4 (refer again to Figures 1.6 and 1.7 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 1.8 we provide a representative example of the results we obtain. Both panels

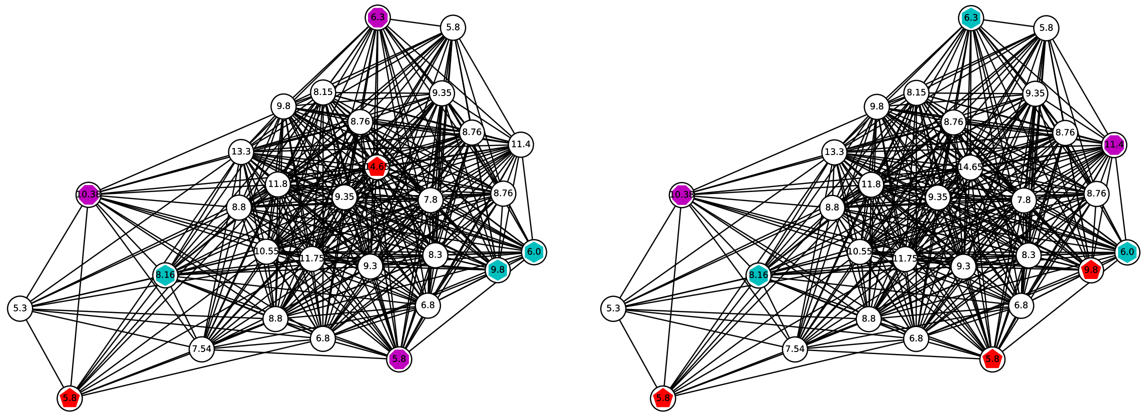

Figure 1.6: **Graph representation of the set of feasible allocations $\mathcal{F}$ of an instance.** Each node represents a trip, labels represent the fare, and two nodes are connected by an edge when they cannot be allocated together. **(left)** A value-maximizing allocation of the jobs to $n = 3$ providers, with the allocation to each provider given by a different color and shape. The total value allocated is \$66.89, and the provider with the smallest allocation receives \$20.45. **(right)** A (value-maximizing) max-min fair allocation. The total value allocated is \$63.64, and the provider with the smallest allocation receives \$20.46.

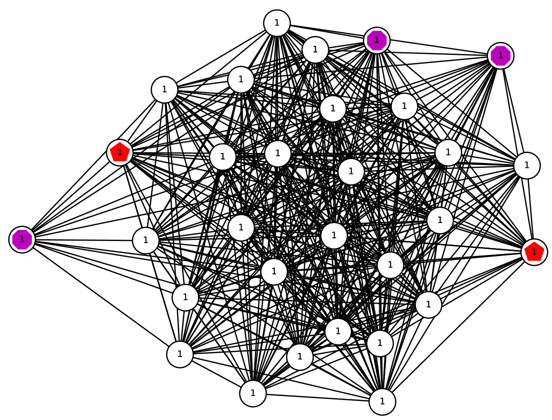

Figure 1.7: **Graph representation of $\mathcal{F}$ for an instance where trips have identical intrinsic values.** Each node represents a trip and two nodes are connected by an edge when they cannot be allocated together. A value-maximizing and Max-Min fair allocation with 2 providers is represented by the coloring and shapes in the nodes. 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.

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.

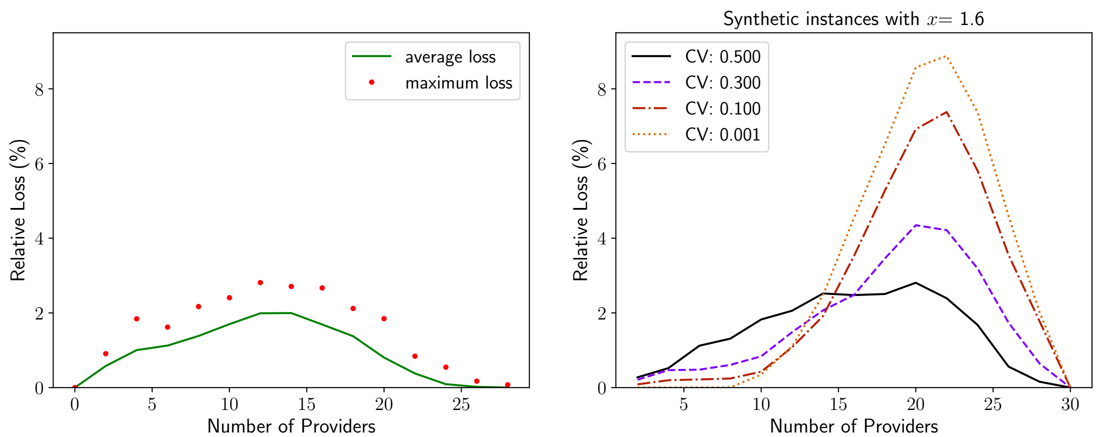

Figure 1.8: **Relative value loss in particular instances. (left)** Average and maximum values of the relative value loss $\mathbf{L}_{\gamma}(\mathcal{F}, \mathcal{F}_{\mathrm{G}})$ as a function of the number of providers, for instances with 30 jobs constructed from the data, using a time window of $w = 16$ minutes. **(right)** Average value of the relative value loss $\mathbf{L}_{\gamma}(\mathcal{F}, \mathcal{F}_{\mathrm{G}})$ as a function of the number of providers for synthetic instances with different coefficients of variation (cv), and using the parametric value $x = 1.6$.

Both charts confirm several of our earlier observations. The right panel in Figure 1.8 shows that the maximum loss decreases with the variability in job values, which is consistent with our discussion in Example 1.1. Additionally, the charts are consistent with the results in Proposition 1.5, and highlight the same qualitative features. For example, the value loss has a unique peak that appears for a ratio of supply to demand between $\frac{2}{3}$ and $\frac{1}{2}$, and decreases to zero when the number of providers becomes sufficiently high or sufficiently low.

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

Theorem 1.1: 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.

1.6 Conclusions, Limitations and 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' effectiveness to generate value when fewer providers are present. We analyzed several additional loss drivers, finding that both a high variation in the intrinsic values of the jobs as well as a very imbalanced (i.e., either very high or very low) ratio of supply to demand would lead to *smaller* losses. Finally, we confirmed several of the findings numerically using both real and synthetic data, and documented that value losses in specific practical settings may be significantly lower than the (theoretical) worst-case values.

These findings and certain limitations in our modeling framework motivate future work of both theoretical and practical nature. Having established that losses are bounded in a general setting and under a broad class of provider guarantees, one could also seek to quantify losses in more specific settings or for subclasses of guarantees obtained as special cases of our framework. For instance, one could seek a parametrized upper-bound on the relative loss when guaranteeing a specific 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 §1.5. These experiments also suggest that it may be worth studying the average-case loss as opposed to the worst-case loss, both due to the substantial gap between the two, but also because certain qualitative effects may be provable under the former setting but not the latter; for instance, the average relative loss appears to be a unimodal function of the number of providers, whereas the worst-case loss does not.

Several of our results also suggested that the degree of symmetry in set of feasible allocations can play a critical and non-trivial role: while asymmetry could in principle lead to unbounded worst-case losses, allowing some asymmetry could actually help reduce empirical losses in some cases of practical interest. To that end, it would be very interesting to understand more deeply the role of symmetry, e.g., by relaxing our standing Assumption 1.1(ii) and introducing a tractable notion that allows a bounded degree of asymmetry in the providers' capability to execute jobs.

Lastly, and from a more prescriptive viewpoint, losses remaining bounded for many types of guarantees also 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.

Chapter 2

Improving Smallholder Welfare While Preserving Natural Forest: Intensification vs. Deforestation

2.1 Introduction

Out of the 1.4 billion people living on less than US$1.25 per day, one billion are smallholder farmers working on land plots smaller than 2 hectares (Rapsomanikis 2015). Yet, at the same time, over 80 per cent of the food consumed in a large part of the developing world is produced by smallholders (IFAD 2013). This puts smallholders at the center of the global efforts to both reduce poverty and increase agricultural production. The latter being ever more important given the rising food demand that is roughly expected to double by 2050 (Tilman et al. 2011). Motivated by these goals and inspired by the Asian Green Revolution of the 60's and 70's (Hazell 2009), governments and NGOs have been actively implementing programs to increase agricultural productivity (see e.g., Djurfeldt et al. 2005, Rashid et al. 2013).

Many of these programs have had some success in increasing the total agricultural productivity. However, in many cases, they have also contributed to the increasing rates of tropical deforestation (see § 2.1.2). Indeed, agricultural expansion has been widely recognized as one of the main causes of tropical deforestation (Geist and Lambin 2002, Kissinger et al. 2012), which in turn is one of the leading causes of anthropogenic Green House Gas emissions (Houghton 2012).

At the same time, the rising temperatures and increased variability of weather events caused by climate change directly affects smallholders (Nelson et al. 2014). Faced with little capital and higher variability in their costs, smallholders turn to land expansion. This makes it even more important to understand the following question: can agricultural *intensification* be achieved while avoiding *deforestation*?

In order to better understand the answer to this question, we present a dynamic model of farmer operations under liquidity constraints and random production costs. We show that the combination of these two factors plays a major role in determining when intensification—defined as any increase in agricultural productivity—will exacerbate or mitigate deforestation. In particular, we show that reducing the cost of intensification can either increase or decrease the deforestation pressure, depending on how large the marginal intensification costs are compared to the variation in production costs.

2.1.1 Main Contribution

We develop a general dynamic model of a farmer's operations, allowing for both productive and land-clearing decisions under liquidity constraints. This allows us to study the effects of changing the farmer's cost structure on their optimal decisions. Our model confirms many previously found results on how farmer welfare increases when reducing production costs (both the average and variability of costs), and when improving access to loans (e.g., Asfaw et al. 2012). Furthermore, we show that although all of these changes would increase the total intensification effort of farmers, they would also lead to a higher rate of deforestation.

We show that, surprisingly, directly reducing the cost of intensification may have different effects on the rate of deforestation: if the intensification cost is low compared to the production cost variability, reducing the intensification cost *reduces* the deforestation pressure. On the other hand, if the intensification cost is high, a reduction would *increase* the pressure. This result helps to shed light on the contradicting empirical evidence linking intensification promotion and deforestation (see §2.1.3).

Central to our results is our consideration of random production costs that are directly proportional to the total productive land. These random shocks occur frequently in practice due to uncertainty when bringing products to market. For instance, transportation costs may be greatly affected by weather conditions when roads are not paved, or labor costs may be higher than expected during harvest season. These risks in the total production costs combine with liquidity constraints to generate a downward pressure on the deforestation rate: faced with a high risk of having to borrow at high interest rates, farmers react to a reduction in the cost of intensification by increasing their rate of production in a smaller productive area.

Our results highlight the importance of considering the specific operational context when designing policy interventions. As has been shown many times in practice, the indiscriminate application of policies can have significant negative consequences. Our model helps policy makers understand the relationship between intensification and deforestation by categorizing farmers and communities of farmers based on their intrinsic characteristics.

2.1.2 Related Literature

The Environmental Science literature has widely documented that agricultural expansion is the dominant driver of illegal deforestation in developing countries (e.g., Carlson et al. 2018, Geist and Lambin 2002). Although there are many proximate causes of deforestation processes, such as economic and institutional factors, the question of whether yield increases is one of these causes is still very much under debate. There is a significant amount of empirical evidence showing how yield increases may lead to both lower rates of deforestation and, on the contrary, higher deforestation rates (see §2.1.3 for a summary of some of these settings). Stemming from these empirical observations, there has been extensive work in the Agricultural Economics literature to shed light on these seemingly contradictory effects (see Angelsen and Kaimowitz 2001 and Angelsen and Kaimowitz 1999 for excellent reviews of these models). Most of the explanations put forth in this body of work can be broadly categorized into three types: labor supply driven (e.g., Maertens et al. 2006, Shively and Pagiola 2004), driven by the elasticity of demand for the agricultural products (e.g., Jayasuriya et al. 2001), or driven by the different types of utility functions of the farmers (e.g., Angelsen et al. 2001). Our work adds to this discussion by considering the role of random production costs paired with liquidity constraints and showing how these two factors play a major role in determining how intensification will affect deforestation.

To develop our model of farmer operations, we use insights from development economics, operations-finance, and sustainable-operations. From the latter, our model generalizes the farmer dynamic model presented in de Zegher et al. (2018), by allowing for dynamic deforestation decisions and generalizing the concave production functions used. From the operations-finance literature, our formulation resembles the models of dynamic production decisions under limited cash inventory, such as Li et al. (2013) and Ning and Sobel (2017). In our model, farmers experience Guassian production cost shocks, that drive them to informally borrow at high rates from specific agents within their community, this is in line with findings from the Development Economics literature (e.g., Collins et al. 2009). Economists have documented that smallholder farmers have limited access to the formal financial system, and rely on informal loans within their communities (Duflo and Banerjee 2011), paying interest rates that increase in the size of the loans. To model these increasing interest payments we adopt an exponential function (see e.g., Ghosh et al. 2000).

In the Environmental Economics and Mechanism Design literature there has been recent interest in designing optimal mechanisms to halt deforestation and preserve natural ecosystems (see, e.g., Mason and Plantinga 2013, Li et al. 2020). Most of these works have been focused on the design of Payments for Ecosystem Services, and not on the detailed operations of farmers. Although in our work we aim at establishing mechanisms for forest protection, we concentrate on the farmers' production operations, and not on the principal-agent problems that arise from the possible conservation mechanisms.

Our work connects as well with the growing body of work in the operations management literature aimed at improving farmers' welfare and social welfare in agricultural supply chains (see Bouchery et al. 2016 and Kalkanci et al. 2019 for reviews). Several recent studies have focused on the production operations of farmers (e.g., Dawande et al. 2013, Boyabatlı et al. 2019, Federgruen et al. 2019, Hu et al. 2019, Levi et al. 2020), as well as the effects of different subsidy schemes on farmer's decisions (e.g., Chintapalli and Tang 2018, Alizamir et al. 2019, Akkaya et al. 2021). In our work we not only examine mechanisms that lead to higher farmer welfare, but study dynamic deforestation decisions.

2.1.3 Examples

In this section we present several empirical examples of how promoting intensification has lead to both increase and decrease of deforestation pressures. The purpose of this section is to illustrate how our result connects to the empirical literature, but it is not an exhaustive review of all such works (see Angelsen and Kaimowitz (2001) and references therein for a broader review).

Governments and NGOs routinely implement two types of programs to incentivize intensification; the first aimed at reducing the cost of inputs, such as fertilizers or pesticides (see, e.g., Pelletier et al. 2020), the second aimed at decreasing the cost of new technology adoptions, such as higher yield seeds or better agricultural practices (see, e.g., Garrett et al. 2013).

In the case of reducing the cost of inputs, this is usually done through either subsidies or low interest rate credits. This is the case for the fertilizer subsidy programs implemented in Zambia and Malawi (see Pelletier et al. 2020, Abman and Carney 2020). Interestingly, although both programs were implemented with similar goals and were indeed successful in increasing yields and farmer welfare, the evidence suggests that in Zambia increased fertilizer use was weakly linked to increased deforestation, while in Malawi, the reverse was found.

In order to reduce the cost of new technology adoption, interventions usually include a combination of training and subsidies for the purchase of new improved inputs (e.g., better seeds). In the case of the Brazilian policy of providing low interest rate credits for the purchase of higher producing soybeans (that was put into place at the end of the 20th century) the results suggest that the improvement in yields led to a higher deforestation rate of the Amazonian forest (Garrett et al. 2013). In Malawi and Zambia, together with the fertilizer programs, the governments implemented high subsidies for the purchase of high-yield maize seeds (Pelletier et al. 2020, Abman and Carney 2020). In contrast to what happened when subsidizing fertilizer, the yield increase caused by the new maize-seeds reduced deforestation in both countries. This same effect was observed in Bangladesh, after government programs subsidized higher yielding crops (Aravindakshan et al. 2021).

Finally, examples of training in better agricultural practices and technologies can be found Indonesia and Malaysia (e.g., Maertens et al. 2006, Villoria et al. 2013). In the case of Indonesia, Maertens et al. (2006) differentiate between the effects observed by yield-saving technologies (that reduced deforestation) and labor-saving technologies (that increased deforestation). While Villoria et al. (2013) observed an increase in deforestation related to higher yields in the oil-palm value chain. Table 2.1 provides a summary of these varied empirical findings.

Country	Type of Intervention	Does Intensification Lead to Deforestation?	Reference
Brazil	Credits for purchasing better yielding soybeans.	Yes,higher yields led to increased deforestation in the absence of strong regulations.	Garrett et al. (2013)
Indonesia (Lore Lindu)	Introduction of labor saving and yield increasing technologies.	Labor saving technologies increased deforestation, yield increasing technologies reduced deforestation.	Maertens et al. (2006)
Indonesia and Malaysia	Training in better agricultural practices and technologies.	Yes, higher yields were associated with higher deforestation.	Villoria et al. (2013)
Zambia	Subsidy of fertilizer and improved maize seeds.	Fertilizer subsidy was weakly linked to increased deforestation; Improved seeds use was linked to a decrease in deforestation.	Pelletier et al. (2020)
Malawi	Subsidies for fertilizer and higher yield seeds.	No, lower rates of deforestation were observed.	Abman and Carney (2020)
Bangladesh	Introduction of higher yielding crops.	No, lower rates of deforestation were observed.	Aravindakshan et al. (2021)

Table 2.1: **Empirical evidence on the Intensification-Deforestation connection.** Summary of some of the many works showing how intensification can either cause or prevent deforestation.

2.2 Model Formulation

We model the operations of a liquidity constrained smallholder-farmer that at every period faces a random production cost shock and must decide on both consumption and production decisions. At the start of each discrete production period n, the farmer observes the current market price p_n, exogenously fixed, and makes three decisions: the rate of consumption c_n, the total amount of land to expand l_n^d, and the total amount of productive inputs or technologies used per unit of time and per unit of land, y_n (henceforth we shall refer to this term as the *productive expenditure rate*). The productive-expenditure may represent the total amount of certain inputs used (e.g, fertilizer, pesticides, insecticides, or labor for preparing the field, planting, and weeding) or the level of adoption of productive technologies (e.g., higher-yield seeds, increased water use), and is characterized by having a concave increasing effect on the total yield.

Formally, consider the time interval $[0, D]$ divided into production periods of length τ. Let the n-th time period be the interval $[(n-1)\tau, n\tau]$ (we assume for simplicity that D is a multiple of τ), and let $N := \frac{D}{\tau}$ be the total number of periods. The timing of decisions is then as follows: at the beginning of the n-th period (i.e., time $(n-1)\tau$), the farmer observes the market price p_n, which we assume comes from an exogenous random process, and that will be paid at the end of the n-th production cycle (i.e., time $n\tau$). Additionally, at the start of the n-th period, the farmer has a cash position x_n, and total productive land ℓ_n. At this time the farmer decides the consumption rate per unit of time c_n, the productive-expenditure rate $y_n \geq 0$, and the total land to expand during period n, ℓ_n^d. The productive-expenditure rate will generate a rate of production given by $(y_n)^\lambda \ell_n$ during period n, for a fixed $0 \leq \lambda \leq 1$. This leads to a total production of $(y_n)^\lambda \ell_n \tau$, during the nth period. Although the land expansion occurs at the start of the period, the land expanded will not become productive until the next period (i.e., time $n\tau$). The total consumption during period n will be $c_n \tau$. At the end of each period, the farmer receives a payment of $(y_n)^\lambda \ell_n \tau p_n$. The land expansion at the start of the period will have a total cost of $(\ell_n^d)^+ d$, where d is the combination of the cost of clearing and making the new land productive. During period n, the farmer will face a total production cost of

$$((y_n)^\lambda \mathcal{W}_n + y_n q)\ell_n \tau. \tag{2.1}$$

Where q is the linear cost of the productive-expenditure y_n (which we will refer to as the *cost of intensification*), and $\mathcal{W}_n$ is a random production cost shock normally distributed with mean k and variance σ^2. Finally, we capture the lack of access to financial markets of the farmer by considering interests $(e^{-\alpha\tau x_n} - 1)$ to be paid at the end of each period. A farmer can only incur debt ($x_n < 0$) by borrowing for one production cycle from within the community, and likewise can lend excess money to the community members. The exponential function is capturing the increasing marginal rates of borrowing (and decreasing marginal rates for lending). This leads to the following dynamics for the farmer's cash position:

$$x_{n+1} = x_n + \underbrace{((y_n)^\lambda \ell_n \tau) p_n}_{\substack{\text{production}\\ \text{revenue}}} - \underbrace{c_n \tau}_{\text{consumption}} - \underbrace{(r(y_n)\mathcal{W}_n + q y_n)\ell_n \tau}_{\substack{\text{production}\\ \text{and intensification}\\ \text{costs}}} - \underbrace{(e^{-\alpha\tau x_n} - 1)}_{\text{interest payments}} - \underbrace{(\ell_n^d)^+ d}_{\substack{\text{deforestation}\\ \text{cost}}} \tag{2.2}$$

Additionally, the land expansion decision ℓ_n^d leads to the following land dynamics:

$$\ell_{n+1} = \ell_n + \ell_n^d. \tag{2.3}$$

Subject to these dynamics the farmer will maximize her expected discounted consumption:

$$\mathbb{E}\left[\int_0^D e^{-\tau\beta} c_{\lceil t/\tau \rceil} dt\right] = \hat{\beta}\mathbb{E}\sum_{n=1}^{N+1} e^{-n\beta\tau} c_n. \tag{2.4}$$

Where $\hat{\beta} = \frac{1-e^{-\beta\tau}}{\beta}$, and the terminal condition is $c_{N+1} = \left(x_{N+1} - (e^{-\alpha\tau x_{N+1}} - 1)\right)/\tau$, which corresponds to the consumption of all the cash remaining net of interest payments. In the objective, β represents the farmer's discount rate. We refer to the farmer's expected discounted consumption as the farmer's *welfare*.

2.2.1 Modelling Assumptions

Timing of farmer's decisions. We assume that farmers decide on their consumption and productive-expenditure rates, as well as the total amount of land they will clear at the start of each period. This assumption not only helps with tractability, but is rooted in practice. Farmers often make production and consumption decisions at the time of cash inflow (see e.g., Collins et al. 2009, Duflo and Banerjee 2011).

Cleared-land production lag. The land cleared at the start of each period will only be considered productive for the next period. This assumption captures the lag between starting to clear land and harvesting from that land. This lag has two main sources: first, clearing land is usually done with manual labor which takes a considerable amount of time (Ketterings et al. 1999), second, once the crops are planted, the time until productive maturity may vary between 100 to 200 days for crops such as maize and ginger (India-Agro-Net 2021a,b) to 42 months for perennial crops such as oil palms (Verheye 2010). Although we are assuming, for simplicity of exposition, the lag to be of one period, all our results can be readily extended to a fixed arbitrary lag T (i.e., $\ell_{n+1} = \ell_n + \ell^d_{(n+1)-T}$).

Production function. We consider the effect on the yield from a production-expenditure of y to be $(y)^\lambda$, for a fixed $0 \leq \lambda \leq 1$. This is in line with the notion of decreasing marginal returns that are found in almost all production technologies. In particular, if we consider y to represent the total rate of labor dedicated to production, by increasing the amount of labor, the farmer can increase the total production, but the rate of increase per unit of labor will be decreasing (Shephard and Färe 1974). If we consider y to represent the rate of fertilizer application, then the function y^λ captures the yield response curve, which has been thoroughly documented to be concave, and is commonly estimated as a power function, with the most common exponents used being $\lambda = 0.5$ and $\lambda = 0.75$ (Tilman et al. 2011, Bélanger et al. 2000, Cerrato and Blackmer 1990, Hagin 1960).

Production cost shocks. We consider random production cost shocks in (2.1) given by $\mathcal{W}_n(y_n)^\lambda \ell_n$. A primary example of these random cost shocks is the delivery cost faced by many smallholders in frontier regions, where roads are seldom paved, leading to highly increased costs when there is enough rain to turn the dirt into mud. Additional random costs associated to bringing the product to market may be linked to higher than expected labor costs at harvesting season.

Market price process. We assume that the market price received by the farmer, p_n, is exogenously determined. This is consistent with many situation in which farmers produce commodity products, such as maize, oil-palm, or cocoa, where the price is mostly fixed by the international markets and not affected by the farmer's own production quantity. These settings are of particular importance, as many of the major documented cases of tropical deforestation are linked to the production of such crops (see Gatto et al. 2017 for a reference on oil-palm production in Indonesia, Bruun et al. 2017 for the case of maize production in the highlands of Thailand, and Kroeger et al. 2017 for an account of deforestation in the Cocoa supply chain). We will assume that $p_n \geq k$, for every n, this assumption avoids the case where production is trivially not sustainable. Additionally, we will assume that $\mathbb{E}(p_{n+1}|\sigma(\{p_i\}_{i\leq n}))$ is increasing in p_n. This implies that observing higher current prices does not lead to lower expected prices in the future. This is consistent with a wide variety of stochastic processes, including any Markovian price process, as well as any submartingale adapted to $\sigma(\{p_i\}_{i\leq n})$.

Negative consumption. While we allow the farmer's consumption c_n to become negative (which can be interpreted as farmers borrowing food from friends and family), we penalize this in the farmer's welfare function (2.4). This assumption is needed for tractability.

Exponential interest payments. Increasing interest rates for larger loans have been extensively documented in the development literature (Duflo and Banerjee 2011, Collins et al. 2009, Ghosh et al. 2000). We capture these increasing loan rates by using an exponential function $e^{-\alpha\tau x_n}$. We assume as well that no farmer would forgo current consumption in order to lend money and use the interest earned in the future (i.e., $\alpha \leq (e^{\beta}-1)/\tau$).

2.3 Results

Theorem 2.1 characterizes the farmer's optimal policy.

Theorem 2.1 *In each period n, the farmer will choose the following productive-expenditure* *and consumption rates, as well as total land cleared:*

$$\ell_n^d = (\hat{\ell}_{n+1} - \ell_n)^+, \tag{2.5}$$

$$y_n = y_n^*(\ell_n, p_n) \tag{2.6}$$

$$c_n = \frac{1}{\tau}\left(x_n + p_n(y_n^*)^\lambda \ell_n \tau - (q y_n^* + k(y_n^*)^\lambda)\ell_n \tau - (e^{-\alpha\tau x_n} - 1) - (\ell_n^{d*})^+ d - g_n^*\right) \tag{2.7}$$

Where $g_n^* = \frac{1}{\alpha\tau}\left((\ell_n(y_n^*)^\lambda \alpha\sigma\tau^2)^2/2 - \log(\frac{e^{\beta\tau}-1}{\alpha\tau})\right)$*, and* $y^*(\ell_n, p_n)$ *solves:*

$$(y^*)^{\lambda-1}\lambda\ell_n\tau(p_n - k) - (y^*)^{2\lambda-1}\lambda\ell_n^2\tau^3\sigma^2\alpha(1 - e^{-\beta\tau}) = q\ell_n\tau. \tag{2.8}$$

A recursive expression for $\hat{\ell}_{n+1}$ can be found in Proposition B.1.

The farmer's best response is divided into three decisions per period, two production decisions (y_n, ℓ_n^d), and one consumption decision (c_n). Interestingly, the production decisions do not depend on the cash position in period n, x_n. This result can be shown to hold for any concave increasing production function and any convex decreasing interest payment function.

The land expansion decision ℓ_n^d follows a base-stock policy form, by which the farmers expand up to a land target $\hat{\ell}_{n+1}$ (and do not expand at all if this target value is below the current amount of land ℓ_n). This target $\hat{\ell}_{n+1}$ is defined in Proposition B.1 in the Appendix as the land amount that equates the expected marginal future value of land to the marginal land expansion cost d. It can be shown that if we assume the lag until new land becomes productive to be T (i.e., $\ell_{n+1} = \ell_n + \ell^d_{(n+1)-T}$), then the optimal land expansion in period n would be $\ell_n^{d*} = (\hat{\ell}_{n+T} - \ell_n - \sum_{i=1}^{T-1} \ell^d_{n-i})$, where the target land $\hat{\ell}_{n+T}$ would equate the marginal deforestation cost d to the expectation in period n of the marginal value of land in all periods following the $(n+T)$-th period.

The optimal production-expenditure rate decision leads to a total production function that increases with the price p_n. Additionally, we show in 2.3 that the optimal production-expenditure decreases with the marginal cost of intensification q and with the variance of the production cost shock, σ^2. This latter relationship can be explained through the high interest payments the farmers face: higher levels of risk will induce lower levels of production intensification. This is consistent with the literature on technological adoption and intensification (see, e.g., Joffre et al. (2018)).

The consumption decisions imply the farmer saves in expectation exactly up to g_n^*, which is increasing in the variance of the production cost shocks. This is consistent with the empirical findings on agricultural risks and how they affect a farmer's ability to obtain food security and higher welfare levels (see e.g., Wolgin 1975). This is even more relevant than ever when facing higher climate-change related risk (Harvey et al. 2014).

Theorem 2.2 shows how the farmer's value function changes with the total amount of land, the intensification cost, the expected cost of production, and the variation in this cost of production.

Theorem 2.2 *The farmer's value function $J_n(x_n, \ell_n, p_n)$ is **increasing** in ℓ_n, p_n, and x_n, and **decreasing** in q, k, and σ^2, for every $n \in \{1, \ldots, N+1\}$.*

As expected, the farmer's welfare is increasing in the total amount of land, the market price, and the cash position, and decreasing in the cost of both the cost of intensification and the expected cost of production. Moreover, the higher the variation of production costs, the lower the total welfare.

In Theorem 2.3 we show how the optimal production-expenditure levels change as a function of the amount of land, the cost of intensification, and the variance of the production cost shocks.

Theorem 2.3 *The farmer's optimal production-expenditure level $y_n^*(\ell_n, p_n)$ is **decreasing** in σ^2, ℓ_n, q, α, and k.*

We show that the production-expenditure exerted is decreasing in the production risk. High variability in production costs generate reduced consumption as well as a reduction in the total optimal production. The more subtle insight we show in Theorem 2.3 is that the production-expenditure level is decreasing in the amount of land: under liquidity constraints, the higher the amount of land, the less farmers can invest in increasing the productivity per unit of land. Increased intensification cost q would, as well, decrease the total production-expenditure rate, which is in line with all the literature on incentivizing technology adoptions and better farming processes, and forms the basis of most of the input-driven incentives and technological training incentives applied widely in practice (see § 2.1.3 for an account of several such incentive programs).

Theorem 2.4 shows how the land-expansion decisions are affected by the expected production cost shock, its associated variance, and the interest rate.

Theorem 2.4 *The farmer's equilibrium deforestation decision* ℓ_n^{d*} *is* ***decreasing*** *in* k, σ^2, *and* α.

This result shows one of the key problems of most incentives schemes with the dual aim of increasing farmer welfare and decreasing land-expansion: most factors that improve the former increase the latter. This phenomenon has been described in the Economics literature as the Jevon's paradox, and indeed occurs frequently in practice (see Alcott 2005). In our model, we can see that decreasing the expected cost of production k or the interest rate α imply an increase in deforestation pressure. Additionally, we can see that at higher levels of variability, the deforestation pressure is reduced. Not only does high variability of costs induce lower intensification, but it reduces as well the amount of land cleared. The rationale for why this happens is similar to before: at higher variability of production costs and faced with high interest rates for debt, farmers are less prone to increase their total productive land.

In Theorem 2.5 we show that for low enough cost q, the land-expansion pressure is actually **increasing** in q.

Theorem 2.5 *There exist positive thresholds* $\tilde{q}_n^L(\ell_n) \leq \tilde{q}_n^H(\ell_n)$ *such that farmer* f*'s equilibrium deforestation decision* ℓ_n^{d*} *is* ***increasing*** *in* q *for* $q \leq \tilde{q}_n^L(\ell_n)$, *and* ***decreasing*** *in* q *for* $q \geq \tilde{q}_n^H(\ell_n)$. *Moreover,* $\tilde{q}_n^L(\ell_n)$ *and* $\tilde{q}_n^H(\ell_n)$ *are increasing in* σ^2, α, *and* ℓ_n.

This surprising result provides a clear insight into the contradicting empirical observations on how intensification can affect deforestation (see §2.1.2 and § 2.1.3). While many subsidy programs that reduced the cost of intensification did reduce the total deforestation, many other have had the exact opposite effect. We demonstrate here that indeed the effect can go in both directions, depending on context specific parameters. This threshold behavior is driven by the combination of the liquidity constraint and the variable production costs. In particular, when reducing the cost of intensification, there are two opposing forces acting on the deforestation pressure. One the one hand, reducing q reduces the intensification cost and increases the equilibrium production-expenditure y^* (see Theorem 2.3), both of which makes each unit of land more valuable. On the other hand, the increased intensification implies that each unit of land will produce a higher volume, which leads to a higher variability in the production costs. Under the liquidity constraints, this increase in total variability will induce a downward pressure on land expansion. The balance between these two forces is characterized by the threshold behavior described in Theorem 2.5: when the variability in production costs is high, $\tilde{q}_n^L(\ell_n)$ will be high, and deforestation pressure will decrease when decreasing q for any q smaller than $\tilde{q}_n^L(\ell_n)$, but when the variability in production costs is low compared to q, q will be above $\tilde{q}_n^H(\ell_n)$, and decreasing q will incentivize deforestation.

Interestingly, Theorem 2.5 shows that the reducing the intensification cost q would not only cause farmers that are "better off" to reduce deforestation. Because $\tilde{q}_n^L$ and $\tilde{q}_n^H$ are increasing in both ℓ_n, and α, then both farmers with an already large productive land, and farmers that are more liquidity constrained would reduce deforestation when their cost of intensification is reduced.

The threshold result in Theorem 2.5, together with the insight on the role that production cost variability and liquidity constraints play, provide an explanation to the highly debated question of whether intensification causes or prevents deforestation (see §2.1.2). To the best of our knowledge this is the first result that presents the level of production cost risk paired with the liquidity constraints as causes for the different answers to this question observed in practice.

2.3.1 Discussion on different incentive schemes

From the results above, we can surmise the following insight.

Insight 2.1 *The only non-conditional welfare improving interventions that can decrease deforestation pressure are those that decrease the intensification cost when this cost is low enough.*

This insight is a direct corollary of Theorems 2.2, 2.4, and 2.5, as any reduction of either the mean of the random production cost or the variance would indeed improve the welfare of the farmer, but would as well lead to an increase in the total land cleared. In contrast, when q, the intensification cost is lower than the threshold $\tilde{q}_n^L$, lowering this cost induces both an increase of welfare and an increase of the protected forest. This insight is validated by the empirical evidence that links the reduction in transportation costs with the increase of deforestation (see Geist and Lambin (2002) for a global analysis of this, and Bruun et al. (2017) for an example of how the improved road conditions facilitated the deforestation of the highlands in northern Thailand). This reenforces the need for careful implementation of policies, because in most cases, well intentioned welfare-improving policies can have devastating effects on the preservation of natural forests if not made conditional on conservation goals.

2.4 Concluding Remarks

We have introduced a dynamic model of farmer operation that allowed us to study the effect of intensification promotion on deforestation. In particular, we found that there exists a non monotonic relationship between the intensification cost and the rate of deforestation: for low enough intensification costs, decreasing the cost can reduce deforestation, while for large enough costs the opposite is true. This adds a nuanced explanation to the already existing theories that try to explain the, empirically observed, varied relationship between intensification and deforestation. In particular, our main result is driven by our consideration of random production cost shocks and liquidity constrained farmers.

Our results reinforce the importance of careful study of each context before any policy implementation. Implementing a policy of intensification promotion in a setting where the average size of plots is small, the variability in costs is low, and the intensification costs are already high, may have negative effects on the forest protection. At the same time, implementing the same policy in a region where there is high variability of costs, land sizes are not too small, and the cost of intensification is not too high may actually reduce deforestation.

Finally, we believe these findings motivate future work that could empirically validate our results. Namely, showing that in settings where the reduction of the cost of intensification led to higher deforestation, the variability in costs was low enough to put the intensification cost above the theoretical threshold. And conversely, that in settings where intensification reduced land-clearing, that the intensification cost was indeed below the threshold. Estimating these thresholds in practice would require a careful collection of farm-level data, in order to understand the sources of uncertainty in the costs, as well as estimating the model's parameters.

Appendix A

Appendices to Chapter 1

A.1 Proofs and Examples

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

We now prove Propositions 1.1 and 1.2.

Proof of Proposition 1.1. If we have that $\mathcal{F}_{\mathrm{G}} = \arg\max_{\boldsymbol{A}\in\mathcal{F}} g(\boldsymbol{A})$, then if $\boldsymbol{A} \in \mathcal{F}_{\mathrm{G}}$, 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 $\boldsymbol{B} \in \arg\max_{\boldsymbol{A}\in\mathcal{F}} g(\boldsymbol{A}) = \mathcal{F}_{\mathrm{G}}$, proving that Assumption 1.2 holds. On the other hand, if we assume that Assumption 1.2 holds for a certain $\mathcal{F}_{\mathrm{G}}$, then let us consider the function $g : \mathcal{F} \to \mathbb{R}$, defined by $g(\boldsymbol{A}) = \mathbb{1}_{\mathcal{F}_{\mathrm{G}}}(\boldsymbol{A})$, that takes the value 1 when $\boldsymbol{A} \in \mathcal{F}_{\mathrm{G}}$, and 0 otherwise. Thus, by definition $\mathcal{F}_{\mathrm{G}} = \arg\max_{\boldsymbol{A}\in\mathcal{F}} g(\boldsymbol{A})$. 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 1.2, $\boldsymbol{B} \in \mathcal{F}_{\mathrm{G}}$, which implies that $g(\boldsymbol{B}) = 1 \geq g(\boldsymbol{A})$. Therefore, $g(\cdot)$ satisfies the condition of Proposition 1.1, which concludes the proof. ■

Proof of Proposition 1.2 It is clear that if we have a union of income guarantees, then Assumption 1.2 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 1.2, 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 1.2, $\boldsymbol{C}\in\mathcal{F}_{\mathrm{G}}$, which concludes the proof. ■

Now we will prove Theorems 1.1 and 1.2. We begin by proving an optimality condition that any allocations $\boldsymbol{A}\in\mathcal{F}^*$and $\boldsymbol{B}\in\mathcal{F}_{\mathrm{G}}^*$ must satisfy.

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

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

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 1.1-(iii) that (B_j',A_{-i}) is a feasible allocation. Hence, by Assumption 1.1-(ii), (B_j',A_{-j}) is a also a feasible allocation and by definition of B_j', we have as well that $B_j'\cap B_i=\emptyset$, for any $i\neq j$, which implies by Assumption 1.1-(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 1.2 on $\mathcal{F}_{\mathrm{G}}$, we must have that $(B_1,\ldots,B_j',\ldots,B_n)\in\mathcal{F}_{\mathrm{G}}$. This implies a contradiction, because $(B_1,\ldots,B_n)\in\mathcal{F}_{\mathrm{G}}^*$, 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 A.1, we now prove Theorem 1.1, that shows an upper bound of $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F},\mathcal{F}_{\mathrm{G}})$.

Proof of Theorem 1.1. Given any $\mathcal{F}$ and $\mathcal{F}_{\mathrm{G}}$, we know that any allocations $(A_1,\ldots,A_n)\in\mathcal{F}^*$ and $(B_1,\ldots,B_n)\in\mathcal{F}_{\mathrm{G}}^*$ must satisfy the conditions imposed by Lemma A.1. Let us then consider the following set of scalar variables:

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

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

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

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

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

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

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

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

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

Equivalently, we can show that $v(B_j) = y_j + \sum_{i=1}^{n} w_{ij}$. Using these two reformulations of $v(A_i)$ and $v(B_j)$, we can write the expression in (1.1) that defines $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}})$ as:

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

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 variables x, y, and w. Because for any $\mathcal{F}$ and $\mathcal{F}_{\mathrm{G}}$, the inequalities of (A.2) 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:

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

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:

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

The optimization problem in (A.5) 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:

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

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

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

Let us define $f(\boldsymbol{\gamma}) : [\gamma_{\min}, \gamma_{\max}]^n \to [0, 1]$, as the optimal value of the dual problem (A.7). Now we can see that each term that appears in the objective of problem (A.7) is a 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 (A.7) 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 A.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 γ that maximizes the solution to problem (A.4) 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 A.2 *The optimal value of the linear program (A.5), when* $\gamma_i = \gamma_{\max}$ *for all* $i \in \{1, \ldots, n_0\}$, *and* $\gamma_i = \gamma_{min}$ *for all* $i \in \{n_0+1, \ldots, n\}$, *for all* $n_0 \in \{1, \ldots, n\}$ *is*

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

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

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

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

In this case, consider the following primal feasible point:

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

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

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

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

$\lambda_{ij} = \dfrac{\gamma_{\min}}{n_0}$ for $i, j \in \{1, \ldots, n_0\}$, $\lambda_{ij} = \dfrac{\gamma_{\min}^2}{\gamma_{\max} n_0}$ for $i \in \{(n_0+1), \ldots, n\}$ and $j \in \{1, \ldots, n_0\}$.

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:

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

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

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

Finally, we take the following dual solution:

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

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

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

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 1.1, we proceed to prove Theorem 1.2. For this, we need only to show that we can asymptotically approximate the upper bound proven in Theorem 1.1.

Proof of Theorem 1.2 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 (1.6)-(1.7). For this, we will generalize Instances 1.1 and 1.2, 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 A.1 *Take $D = \{d_1\}$, such that $v(d_1) = 1$, any n, and $\gamma_1 = \gamma_{\max}$, $\gamma_i = \gamma_{\min} \leq \gamma_{\max}$ for all $i \in \{2, \ldots, n\}$. 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 1.2 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 A.2 *Given any n, let $D = \{d_1, \ldots, d_n, \ldots, d_{2n-1}\}$, such that $v(d_i) = 1$, for all $i \in \{1, \ldots, n\}$, and $v(d_j) = 1 - \kappa$, for $j \in \{n+1, \ldots, 2n-1\}$. Let $\gamma_1 = \gamma_{\max}$ and $\gamma_i = \gamma_{\min}$, 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 (1.3) 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 (1.4).

Given this instance, the only efficient allocation would be $\boldsymbol{A}$ with $A_1 = \{d_1, \ldots, d_n\}$, $A_i = \{d_{n+i-1}\}$ for $2 \leq i \leq n$. That is, we assign all of the first n jobs to one provider and distribute the remaining jobs between the remaining providers. Moreover, the only allocation in $\mathcal{F}_{\mathrm{G}}^{\mathrm{mM}}$ (modulo symmetries) is given by $\boldsymbol{B}$ with $B_i = \{d_i\}$, 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 loss $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) =$

$\frac{(n\gamma_{\max}+\gamma_{\min}(n-1)(1-\kappa))-(\gamma_{\max}+(n-1)\gamma_{\min})}{n\gamma_{\max}+\gamma_{\min}(n-1)(1-\kappa)} = \frac{(n-1)\gamma_{\max}-\kappa\gamma_{\min}(n-1)}{n\gamma_{\max}+\gamma_{\min}(n-1)(1-\kappa)} \underset{\kappa\to 0}{\to} \frac{n-1}{n+(1-\delta)(n-1)}$. *Therefore, for any* $\epsilon > 0$, *there exists a* κ *small enough such that* $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) = \frac{n-1}{n+(1-\delta)(n-1)} - \epsilon$.

Instance A.3 *Given any* n, *and* t, *positive integers, let* $D = \cup_{k=1}^{t} C^k \cup S$, *where* $C^k = \{d_2^k, \ldots, d_n^k\}$, *and* $S = \{d_1^s, \ldots, d_{n(t+1)}^s\}$. *Let as well* $\gamma_1 = \gamma_{\max} \geq \gamma_{\min} = \gamma_i$, *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 (1.3) 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 (1.4).*

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

To better visualize the structure of Instances A.2 and A.3, we can imagine a graph on the elements of D, where a set of jobs can be assigned together only when there is no edge between any pair of the corresponding vertices. Thus, in the case of Instance A.2, there

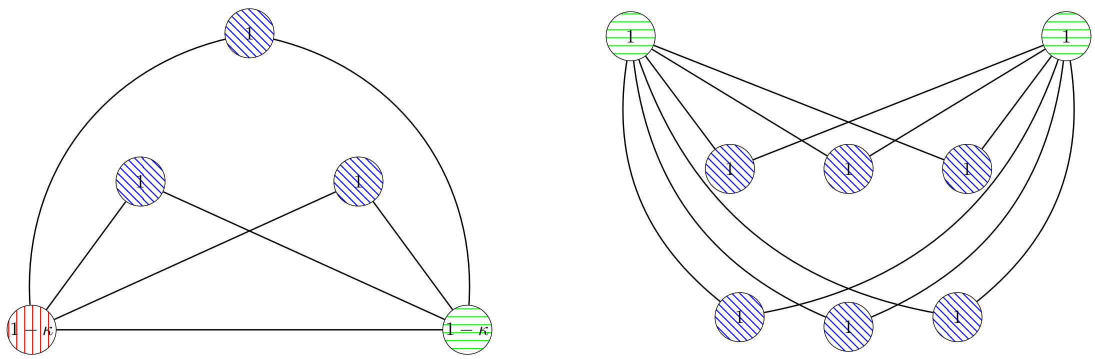

Figure A.1: **Feasibility graphs for Instance A.2 (left)**, when n is 3, and Instance A.3 **(right)**, when n is 3 and t is 1.

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 A.1. These types of graphs are known as complete split graphs (see Le and Peng (2015)). Similarly, in the case of Instance A.3, there would be no edges between any pair of vertices in the set S, 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 A.1. An allocation of the jobs can be seen as a covering of this graph by independent sets (sets of vertices without any edge joining two vertices of the set).

The family of Instances A.1, A.2, and A.3 prove that $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}})$ 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 1.2. ■

We now present Instance A.4, that shows how when we relax Assumption 1.1 (ii), we can achieve a loss that grows asymptotically to 100% with the number of providers.

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

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

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* $p_i(A_i) = \gamma_i \sum_{d \in A_i} v(d)$*. Then*

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

Where $\mathcal{F}_{\mathrm{G}}^{\mathrm{mM}}$ *is the allocations that satisfy Max-Min fairness, as defined in* (1.4)*.*

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 1.1 (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 1.7 to 1.6 from §1.4.

Proof of Proposition 1.7. In order to prove this proposition, we will first formally define the set of guarantees $\mathcal{F}^{c}_{G}$, as any subset of $\mathcal{F}^{c}$, that satisfies Assumption 1.2, replacing in the definition of the assumption $\mathcal{F}$ by $\mathcal{F}^{c}$, and using the generalized notion of $v(C_i)$ for $C \in \mathcal{F}^{c}$, described in §1.4. In particular, we extend the notion of uniform income guarantees under monotonic payment functions: given $p(\cdot)$, a monotonic payment function, we take $\mathcal{F}^{c}_{G} = \{\boldsymbol{C} \in \mathcal{F}^{c} \mid p(C_i) \geq \tau, \text{ for } i \in N\}$. As in §1.2.3, it is easy to see that these guarantees satisfy the extended version of Assumption 1.2. Now, we will prove that given any uniform income guarantee, $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}^{c}, \mathcal{F}^{c}_{G}) = 0$.

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 $\boldsymbol{C} \in \mathcal{F}^{c}$, 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 $\mathcal{F}^{c}_{G}$ such that $\boldsymbol{C} \notin \mathcal{F}^{c}_{G}$. 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, $\boldsymbol{C} \notin \mathcal{F}^{c}_{G}$ 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 $\boldsymbol{B} \in \mathcal{F}^{c}$ 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 1.3.

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 1.2, 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 1.2, 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 1.2. ■

Proof of Proposition 1.4.

We begin by proving (i). Without loss of generality, we will normalize $v(d) = 1$, for all $d \in D$. To show that we have zero loss under any uniform income guarantee $\mathcal{F}_{\mathrm{G}}$, 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 1.5 is a generalized form of Instance 1.2, 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 1.2, when $n = 2$, and Instance A.2 when $n \geq 2$. ■

We now describe Example A.1, that shows how the dependency of the loss on the variance of the values described in Example 1.1 extends to the n provider case.

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

$$\begin{aligned}
\mathbb{E}_\kappa(\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}})) &= \int_0^{\min\{1, \frac{\Delta}{2}\}} \frac{1}{\Delta} \frac{(n-1)(1-\kappa)}{n + (n-1)(1-\kappa)}\, d\kappa \\
&= \frac{1}{\Delta}\Big(\frac{n}{n-1} \log((n-1)\min\{1, \frac{\Delta}{2}\} - (2n-1)) \\
&\qquad + \min\{1, \frac{\Delta}{2}\} - \frac{n}{n-1} \log(2n-1) \Big) \\
&:= g(\Delta).
\end{aligned}$$

As in Example 1.1, 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 1.5. We will first prove that if $n = 1$ and $\gamma_{\min} = \gamma_{\max}$, then $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) = 0$, for any set of allocations with guarantees, $\mathcal{F}_{\mathrm{G}}$. To show this, we simply observe that due to Assumption 1.2, 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 the loss is zero when $n \geq |D|$, for any set of allocations with guarantees. For this, we assume without loss of generality that it is always feasible to allocate at least one job to any specific provider (if not, then there is a job that cannot be completed by any provider, and we could then simply ignore it). Hence, we claim that $\max_{\boldsymbol{A} \in \mathcal{F}} \sum_i v(A_i) = \sum_{d \in D} v(d)$. 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 $j \neq i$, and $A'_i = \{d\}$. This leads to a contradiction, because by Assumption 1.2, $\boldsymbol{A}' \in \mathcal{F}_{\mathrm{G}}$, 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 $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) = 0$, for any set of allocations with guarantees $\mathcal{F}_{\mathrm{G}}$. ■

Proof of Proposition 1.6. Without loss of generality, we can assume that all intrinsic values are 1, that is, $v(d) = 1$, for each $d \in D$. Now, because $|D| < 2n$, and because we can always allocate only one job to any provider, then $\max_{\boldsymbol{A} \in \mathcal{F}} \min_{i \in N} v(A_i) = 1$. 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 $\mathbf{L}_{\boldsymbol{\gamma}}(\mathcal{F}, \mathcal{F}_{\mathrm{G}}) = 0$, under any uniform income-guarantee under monotonic payment functions. ■

A.2 Numerical Analysis of Synthetic and Real-world Data

In this section we provide the details of the numerical analysis we discuss in §1.5. 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, in order to better conform to part (ii) of Assumption 1.1, we took two different approaches. First, we filtered the trips that started and ended in a limited region of Manhattan (we took Midtown, Upper West Side and Upper East Side). By considering this small geographical region, we limit the effect of spacial considerations, and better conform to part (ii) of Assumption 1.1, that any feasible set of jobs could be performed by any provider. For our second approach, although we did not limit the starting and ending region, we added geographical constraints to our allocations that we describe below. By adding these spacial considerations we ensure that each set of feasible trips is geographically consistent, while at the same time we satisfy part (ii) of Assumption 1.1 by allowing any provider to perform any set of feasible trips.

In Figure A.2 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, thus removing several outliers that were clearly due to corrupted data (trips of almost 24 hours or more than $1000).

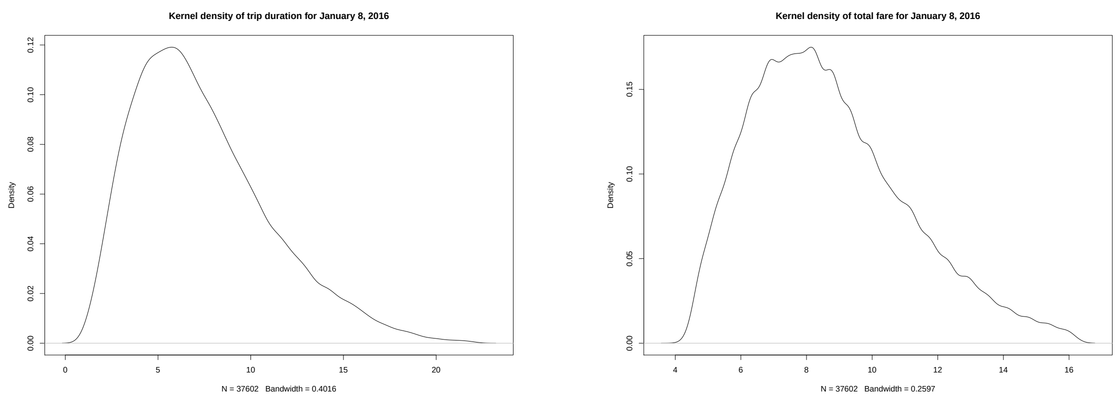

Figure A.2: **(left)** Empirical distribution of the duration of trips. **(right)** Empirical distribution of the total fare of trips.

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 in two ways. For our first approach we only required that trips that intersected in time were not allocated to the same provider. For our second approach, we added the restriction that two trips can only be allocated together if, driving at an average speed of v mph, a driver can reach the start-point of one trip from the end-point of the other trip. We considered speeds of v in the range $\{7.44\text{mph}, 10.7\text{mph}\}$, which were the mean and median speeds across the whole city in 2016 (see New York City Department of Transportation 2018). Then, we solved for both the total value-maximizing solution and for the value-maximizing solution among the Max-Min fair solutions, for a varying number of providers. For this, we used 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 we obtained the two value-maximizing solutions, we computed the total relative loss across the whole time horizon by taking the relative difference of total value with and without the Max-Min fair restriction. We sampled the instances 100 times, and computed for each number of providers the average value loss across these 100 samples.

In Figures A.3-A.4 we can see the relative value loss as a function of the number of providers, without any geographical restrictions, for different combinations of dates, regions, and size of interval, w. In Figure A.5 we can see the relative value loss as a function of the number of providers, under the feasibility constraints mentioned above, for the whole city of New York, for different average speeds v. We show here a representative set of our results, the complete set of results is available upon request.

We can see in Figure A.3 that there does not seem to be much difference from region to region. We observed similar results across the three regions for all the combinations of dates and values of w we tested. In Figure A.4 we see that for the same region but different dates the relative losses do not appear to change much. Nevertheless, by comparing Figure A.3 to Figure A.4 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 A.4 and A.3 we observed the same pattern mentioned in §1.5 regarding the effect of the providers to job ratio on the loss, namely, for extreme values of this ratio the loss collapses to zero, and the loss achieves its maximum value at an intermediate value. Moreover, for every day and region we analyzed we observe that the curves of maximum losses and average losses are relatively close together, implying that the low average losses are due to frequent low losses, as opposed to infrequent high losses.

In Figures A.5 we can see that adding a geographic feasibility constraint does not seem to change significantly the average losses. Furthermore, changing the average speed from the mean bus speed in the city to the median bus speed does not result in any apparent change in the average losses. It is worth noting that increasing the speed v did lead eventually to almost the same results seen in Figure A.4, while reducing v to zero lead to a constant zero loss.

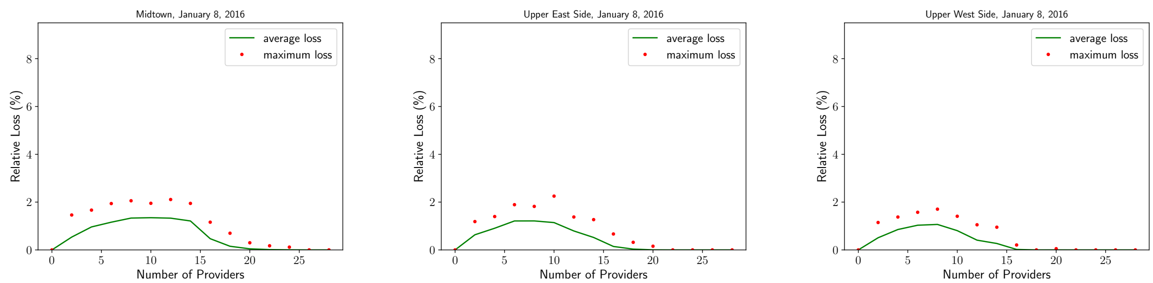

Figure A.3: **Relative value loss for different regions.** Average and maximum $\mathbf{L}_{\gamma}(\mathcal{F}, \mathcal{F}_{\mathrm{G}})$ as a function of the number of providers, for instances with 30 jobs constructed from the data, using $w = 20$ and three different regions of NYC.

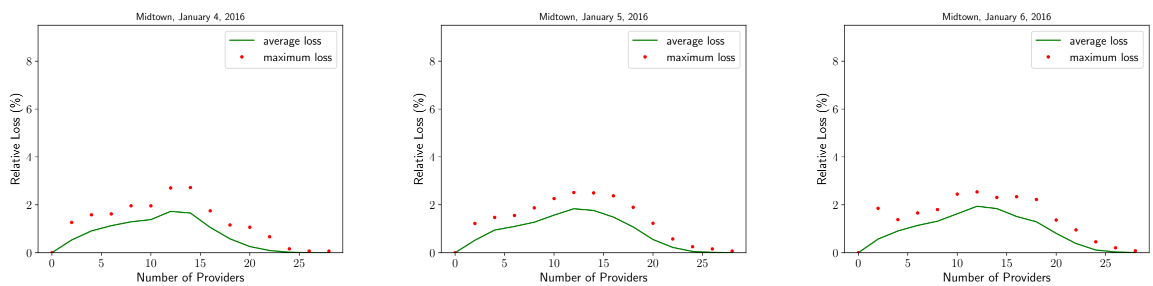

Figure A.4: **Relative value loss for different dates.** Average and maximum $\mathbf{L}_{\gamma}(\mathcal{F}, \mathcal{F}_{\mathrm{G}})$ as a function of the number of providers, for instances with 30 jobs constructed from the data, using $w = 15$ and three different dates of the first week of 2016.

As a robustness test of how our results would change when Assumption 1.1(ii) is relaxed, we also computed the average loss in instances that violate the symmetry assumption. Namely, we parametrized the symmetry of each instance by s: in an instance with symmetry $s\%$, each provider can only perform $s\%$ of all the jobs. To construct instances with $s\%$ symmetry, we took the same set of jobs we considered in our first approach defined above and randomly (and independently) selected the set of jobs each provider can perform. Hence, at 100% symmetry, our instances satisfy Assumption 1.1(ii) and are the

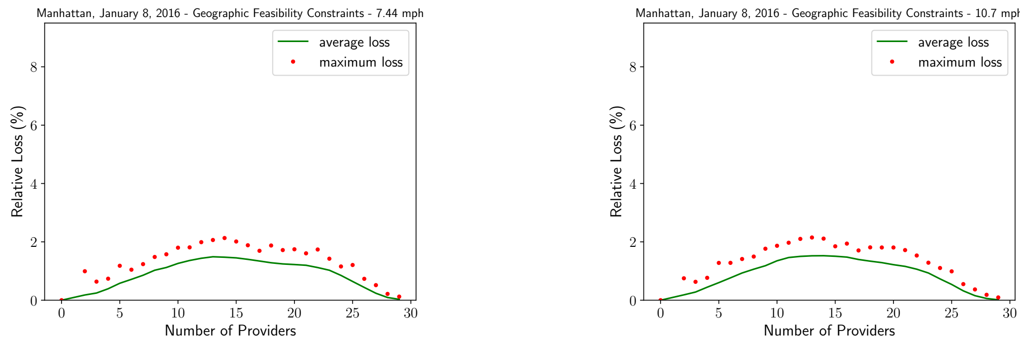

Figure A.5: **Relative value loss with geographic constraints for different average speed, v.** Average and maximum $\mathbf{L}_\gamma(\mathcal{F}, \mathcal{F}_G)$ as a function of the number of providers, for instances with 30 jobs constructed from the data with geographic constraints on the feasible allocations, using $w = 15$ and the two values of v: $v = 7.44$ (**left**) and $v = 10.44$ (**right**)).

same as those analyzed above. In Figure A.6 we can see the average relative loss as a function of the number of providers, for different values of s (with $w = 15$, the region as Midtown Manhattan, and no geographical restrictions). The results suggest that for "high symmetry" (i.e., values of s above 60), the resulting average losses are not significantly different from those under full symmetry, recorded in Figures A.3-A.4; but once instances become "sufficiently asymmetric" (the symmetry parameter drops below 50), the average losses decrease. Intuitively, lower values of s imply lower probabilities of two providers being able to perform the same job, which may explain why for low number of providers and low values of s the loss is always zero: if no job can be performed by two providers, then the allocation problems can be separated into disjoint problems for each provider, which, by Proposition 1.5 implies that the value loss will be zero. This shows that asymmetry can play a nuanced role, depending on other problem parameters: some asymmetry actually reduced the average losses here, but complete asymmetry also lead to worst-case losses that asymptotically approached 100%, as seen in Instance A.4.

Synthetically generated Instances. In order to analyze the dependency between the variation of values and the relative value loss we mention in §1.4, we generate synthetic instances where we can control this variation. In particular, for each instance we sample uniformly 30 starting points in the interval $(0, x] \subseteq \mathbb{R}$, for different x in the interval $(1, 3)$, and for each point we sample from a truncated normal distribution the length of the interval. We take this truncated normal distribution with mean 1 and a coefficient of variation cv varying from 0.001 to 0.5. Each interval represents a specific job (similarly to the trips in the 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 A.7. 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 A.7, as we mentioned in §1.5, that the maximum loss in decreasing in the coefficient of variation of the values, consistent with the remarks of §1.4 on the variation of values as a driver of loss. Nevertheless, in the instances we generated for Figure A.7, the variation of the values is intrinsically connected to the variation in the lengths of the intervals we took to generate the feasibility restrictions. Hence, in order to isolate the effect of the variation of values, we took the same instances, but where we fixed the length of each interval (representing a job) to be exactly 1, for feasibility purposes. Therefore, while the intervals that define the feasibility constraints all have length 1, the values remain as before, taken from truncated normal distributions with different coefficients of variation. We then plotted the average losses as before in Figure A.8. By comparing Figure A.7 and A.8, 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. Moreover it demonstrates the effect of many of the main drivers of loss we analyzed in §1.4.

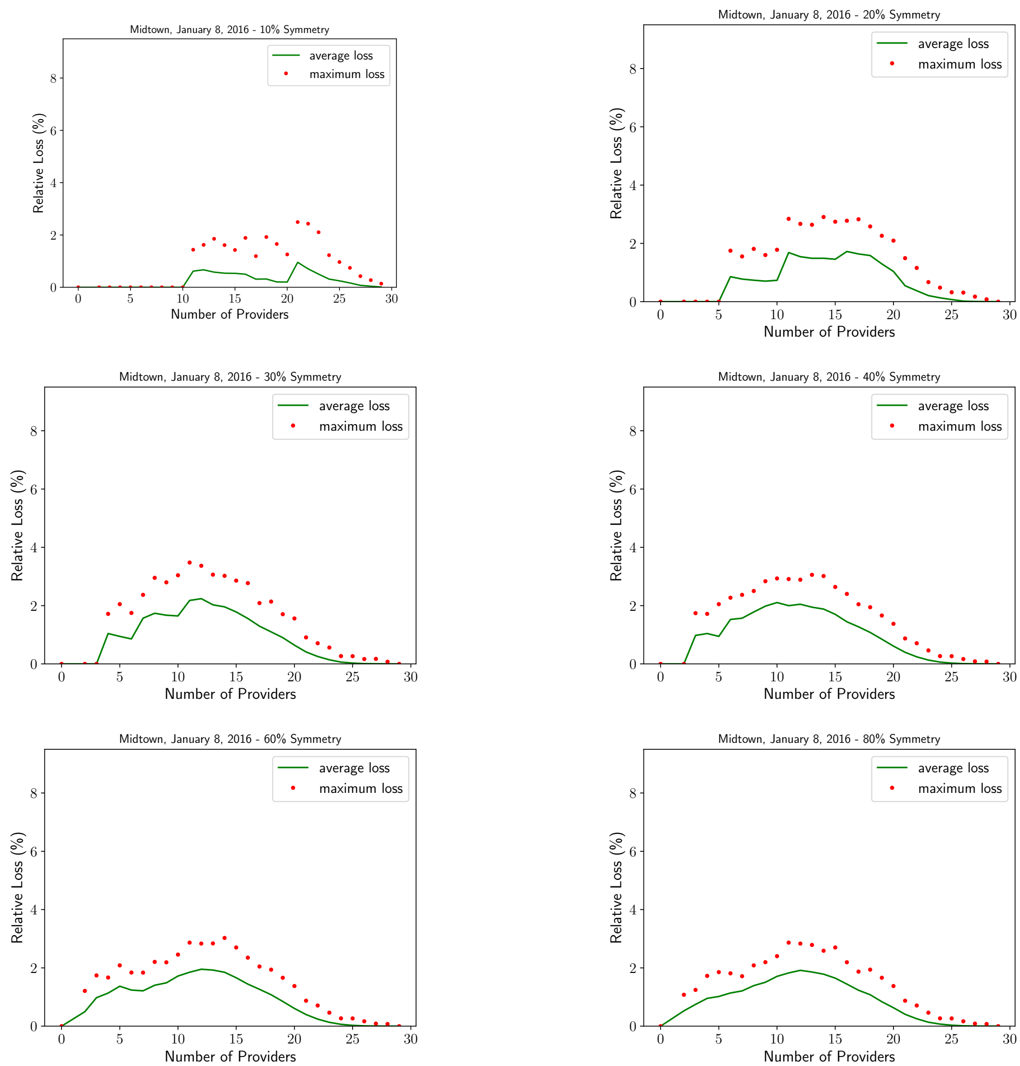

Figure A.6: **Relative value loss for different values of the symmetry parameter, s.** Average and maximum $\mathbf{L}_\gamma(\mathcal{F}, \mathcal{F}_\mathrm{G})$ as a function of the number of providers, for instances with 30 jobs constructed from the data with varying levels of symmetry s, using $w = 15$ and the four values of s: $s = 20\%$ (**upper left**), $s = 40\%$ (**upper right**), $s = 60\%$ (**lower left**), $s = 80\%$ (**lower right**).

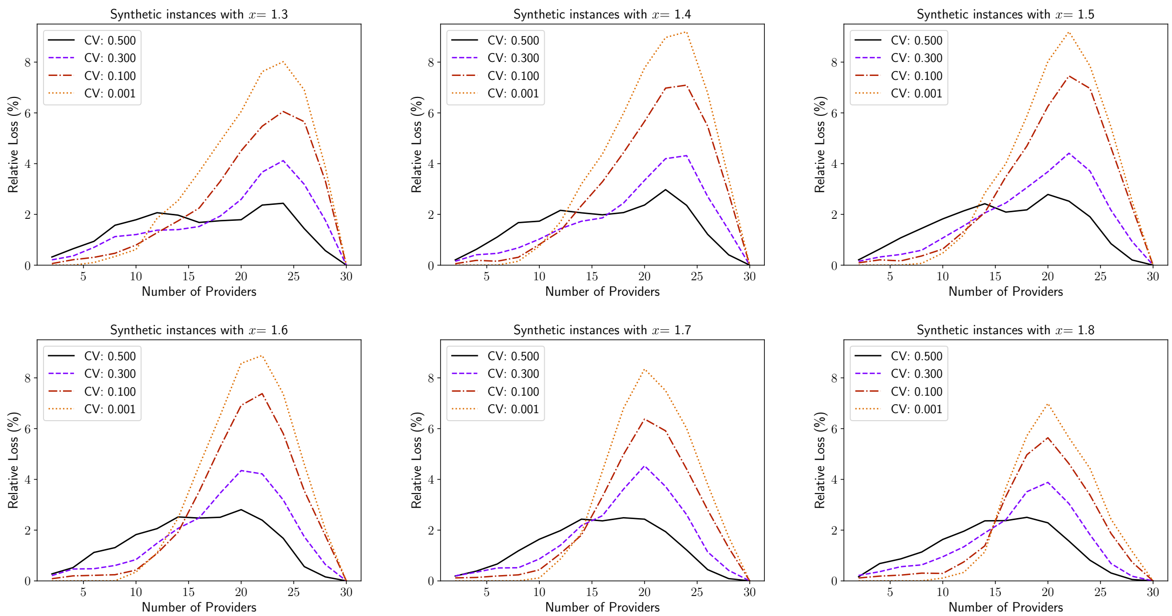

Figure A.7: **Synthetic instances for x varying from 1.3 to 1.8.**

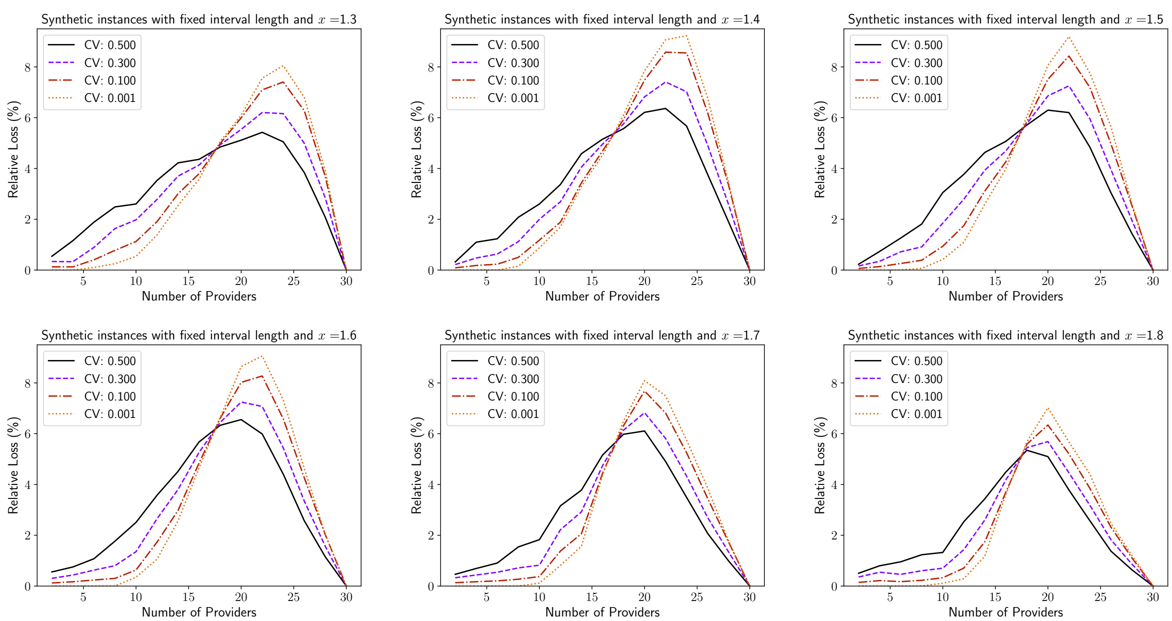

Figure A.8: **Synthetic instances with fixed interval length for x varying from 1.3 to 1.8.**

Appendix B

Appendices to Chapter 2

B.1 Proofs

We begin by proving a proof of Theorem 2.1.

Proof of Theorem 2.1. Recall that the problem is divided into $N+1$ periods of length τ, where the n-th period corresponds to the time $[(n-1)\tau, n\tau)$. In the terminal period, $N+1$, the farmer consumes the remaining cash position net of interest payments. In the rest of the periods, the farmer observes the price p_n, total productive land ℓ_n, and current cash position x_n, and decides on the consumption rate c_n, the production-expenditure rate y_n, and the total deforestation amount ℓ_n^d. Let $J_n(x_n, \ell_n, p_n)$ denote the farmer's value function at time $n \in \{1, \ldots, N+1\}$. We show the following Proposition that proves the desired result:

Proposition B.1 *For $n \in \{1, \ldots, N+1\}$,*

$$J_n(x_n, \ell_n, p_n) = \frac{\hat{\beta}}{\tau}(x_n + (1 - e^{-\alpha x_n}) + f_n(p_n, \ell_n)), \tag{B.1}$$

$$\ell_n^{df} = (\hat{\ell}_{n+1} - \ell_n)^+, \text{ where } \hat{\ell}_{n+1} \text{ solves } \frac{e^{-\beta\tau}\partial \mathbb{E}_n f_{n+1}(p_{n+1}, l)}{\partial l} = d, \tag{B.2}$$

$$y_n(p_n) = y^*(\ell_n, p_n), \tag{B.3}$$

$$c_n^f = \frac{1}{\tau}\Big(x_n + p_n(y_n^*)^\lambda \ell_n \tau - (q y_n^* + k(y_n^*)^\lambda)\ell_n \tau - (e^{-\alpha\tau x_n} - 1) - (\ell_n^d)^+ d - g_n\Big), \tag{B.4}$$

where $g_n = \frac{1}{\alpha\tau}\Big(((\ell_n)^2 (y_n)^{2\lambda}\alpha^2\sigma^2\tau^4)/2 - \log(\frac{e^{\beta\tau}-1}{\alpha\tau})\Big)$, $\hat{\beta} = \frac{1-e^{-\beta\tau}}{\beta}$, *$f_n(\ell_n, p_n)$ is concave and increasing in ℓ_n and increasing in p_n, the expectation $\mathbb{E}_n$ is taken conditional on the σ-algebra $\sigma(\{p_i\}_{i\le n}, \{\mathcal{W}_i\}_{i<n})$, and $y^*(\ell_n, p_n)$ solves:*

$$(y^*)^{\lambda-1}\lambda\ell_n\tau(p_n - k) - (y^*)^{2\lambda-1}\lambda\ell_n^2\tau^3\sigma^2\alpha(1 - e^{-\beta\tau}) = q\ell_n\tau. \tag{B.5}$$

Proof of Proposition B.1. We first show that for $n = N+1$, equation (B.1) holds. In this period, the farmer no longer produces, and consumes at a constant rate $c_{N+1} = (x_{N+1} - (e^{-\alpha\tau x_{N+1}} - 1))/\tau$, that leads to a value function that can be written as:

$$\begin{aligned} J_{N+1}(x_{N+1}, \ell_{N+1}, p_{N+1}) &= \int_0^\tau \left(\frac{x_{N+1} - (e^{\alpha\tau x_{N+1}} - 1)}{\tau}\right) e^{-\beta s} ds \\ &= \frac{1 - e^{-\beta\tau}}{\beta\tau}\left(x_{N+1} - (e^{-\beta\tau x_{N+1}} - 1)\right), \end{aligned}$$

that is consistent with equation (B.1), taking $f_{N+1}(p_{N+1}, \ell_{N+1}) = 0$, which is constant and thus concave increasing in ℓ_{N+1} and increasing in p_{N+1}.

We now proceed by induction in n, we assume the induction hypothesis for $n+1$, and consider the farmer's decision problem at time n. The farmer's value to go function at time n is given by:

$$J_n(x_n, \ell_n, p_n) = \max_{y_n \geq 0, c_n, \ell_n^d \geq 0} \left\{ c_n \int_0^\tau e^{-\beta s} ds + e^{\beta\tau} \mathbb{E}_n \left[J_{n+1}(x_{n+1}, p_{n+1}(\mathcal{P}_{n+1}), \ell_{n+1}) \right] \right\}$$
$$= \max_{y_n \geq 0, c_n, \ell_n^d \geq 0} \frac{\hat{\beta}}{\tau} \left\{ c_n\tau + e^{-\beta\tau} \mathbb{E}_n \left[x_{n+1} - (e^{-\alpha x_{n+1}} - 1) + f_{n+1}(\ell_{n+1}, p_{n+1}) \right] \right\}$$

Where the second inequality is due to the inductive hypothesis. Now, we will define the following auxiliary variable:

$$g_n = x_n + p_n y_n^\lambda \ell_n \tau - c_n \tau - (q y_n + k y_n^\lambda) \ell_n \tau - (e^{-\alpha\tau x_n} - 1) - (\ell_n^d)^+ d,$$

which, using the cash dynamics for the farmer imply that

$$c_n = \frac{1}{\tau}(x_n + p_n y_n^\lambda \ell_n \tau - g_n - (q y_n + k r_n) \ell_n \tau - (e^{-\alpha\tau x_n} - 1) - (\ell_n^d)^+ d),$$

$$x_{n+1} = g_n - y_n^\lambda \ell_n \tau \sigma \varepsilon_n,$$

with $\varepsilon \sim N(0,1)$. Using these identities, we can rewrite the value to go function at time n as

$$= \max_{y_n \geq 0, g_n, \ell_n^d \geq 0} \frac{\hat{\beta}}{\tau} \{ \quad (x_n + p_n y_n^\lambda \ell_n \tau - g_n - (q y_n + k y_n^\lambda) \ell_n \tau - (e^{-\alpha\tau x_n} - 1) - (\ell_n^d)^+ d)$$
$$+ e^{-\beta\tau} \mathbb{E}_n \left[x_{n+1} - (e^{-\alpha\tau x_{n+1}} - 1) + f_{n+1}(\ell_n + \ell_n^d, p_{n+1}) \right] \}$$
$$= \max_{y_n \geq 0, g_n, \ell_n^d \geq 0} \frac{\hat{\beta}}{\tau} \{ \quad (x_n + p_n y_n^\lambda \ell_n \tau - g_n - (q y_n + k y_n^\lambda) \ell_n \tau - (e^{-\alpha\tau x_n} - 1) - (\ell_n^d)^+ d)$$
$$+ e^{-\beta\tau} \mathbb{E}_n \left[g_n - y_n^\lambda \ell_n \tau \sigma \varepsilon_n - (e^{-\alpha\tau(g_n - y_n^\lambda \ell_n \tau \sigma \varepsilon_n)} - 1) + f_{n+1}(\ell_n + \ell_n^d, p_{n+1}) \right] \}$$
$$= \frac{\hat{\beta}}{\tau} \{ (x_n - (e^{-\alpha\tau x_n} - 1) + \max_{y_n \geq 0, g_n, \ell_n^d \geq 0} h(y_n, \ell_n^d, g_n) \}$$

Where

$$\begin{aligned}
h(y_n,\ell_n^d,g_n) &= y_n^\lambda \ell_n \tau p_n - g_n - (qy_n + ky_n^\lambda)\ell_n\tau - (\ell_n^d)^+ d \\
&\qquad + e^{-\beta\tau}\mathbb{E}_n\left[g_n - (e^{-\alpha\tau(g_n - y_n^\lambda \ell_n \tau\sigma\varepsilon_n)} - 1) + f_{n+1}(\ell_n + \ell_n^d, p_{n+1})\right]. \\
&= y_n^\lambda \ell_n \tau p_n - g_n - (qy_n + ky_n^\lambda)\ell_n\tau - (\ell_n^d)^+ d \\
&\qquad + e^{-\beta\tau}(g_n - (e^{-\alpha\tau g_n + (\alpha y_n^\lambda \ell_n \tau^2\sigma)^2/2} - 1) + \mathbb{E}_n f_{n+1}(\ell_n + \ell_n^d, p_{n+1})).
\end{aligned}$$

This last equality follows from taking the Gaussian Moment Generating Function (recall that $\varepsilon \sim N(0,1)$):

$$\mathbb{E}_n\left[(e^{-\alpha\tau(g_n - r_n\ell_n\tau\sigma\varepsilon_n)}\right] = e^{-\alpha\tau g_n + (\alpha y_n^\lambda \ell_n\tau^2\sigma)^2/2}$$

We can separate $h(y_n,\ell_n^d,g_n)$ into two functions,

$$h(y_n,\ell_n^d,g_n,\ell_n,p_n) = h^1(y_n,g_n,\ell_n,p_n) + h^2(\ell_n^d,\ell_n,p_n),$$

where

$$h^1(y_n,g_n,\ell_n,p_n) = \ell_n\tau(y_n^\lambda(p_n-k) - y_n q) - (1-e^{-\beta\tau})g_n - e^{-\beta\tau}(e^{-\alpha\tau g_n + (\alpha y_n^\lambda \ell_n\tau^2\sigma)^2/2} - 1),$$

$$h^2(\ell_n^d,\ell_n,p_n) = \mathbb{E}_n f_{n+1}(\ell_n + \ell_n^d, p_{n+1}) - (\ell_n^d)^+ d.$$

By inductive hypothesis, we know that $f_{n+1}(\ell, p_{n+1})$ is concave and increasing in it's first argument, which implies that $h^2(\ell_n^d)$ is concave in ℓ_n^d. Hence, we can take the first order conditions to maximize $h^2(\ell_n^d)$:

$$\frac{\partial h^2(\ell_n^{d*},\ell_n,p_n)}{\partial \ell_n^d} = 0 \Leftrightarrow \frac{e^{-\beta\tau}\partial\mathbb{E}_n f_{n+1}(\ell_n + \ell_n^{d*}, p_{n+1})}{\partial \ell_n^d} = \mathbb{1}_{\{\ell_n^{d*}\geq 0\}} d. \tag{B.6}$$

From equation (B.6) we conclude that the optimal ℓ_n^d must satisfy $\ell_n^d = (\hat{\ell}_{n+1} - \ell_n)^+$, where $\hat{\ell}_{n+1}$ solves $\frac{e^{-\beta\tau}\partial\mathbb{E}_n f_{n+1}(\ell,p_{n+1})}{\partial\ell} = d$. This proves equation (B.2).

In order to optimize $h^2(y_n, g_n, \ell_n, g_n)$, we will begin by finding $g_n^*(y_n)$. Notice that $h^2(y_n, g_n)$ is concave in g_n, because it is an affine function of g_n minus a convex function of g_n. Hence, we can take the first order conditions on $h^1(y_n, g_n, \ell_n, g_n)$, with respect to g_n, to obtain:

$$\frac{\partial h^1(y_n, g_n^*, \ell_n, p_n)}{\partial g_n} = 0 \Leftrightarrow exp(-\alpha\tau g_n^* + (\alpha y_n^\lambda \ell_n \tau^2 \sigma)^2/2) = \frac{e^{\beta\tau}}{\tau\alpha}. \tag{B.7}$$

By using equation (B.7), we obtain $g_n^* = \frac{1}{\alpha\tau}\left((\ell_n^2 y_n^{2\lambda}\alpha^2\sigma^2\tau^4)/2 - \log(\frac{e^{\beta\tau}-1}{\alpha\tau})\right)$, proving (B.4). Now, we can write $h^3(y_n, \ell_n, p_n) = \max_{g_n} h^1(y_n, g_n^*(y_n), \ell_n, p_n)$, as:

$$\begin{aligned} h^3(y_n, \ell_n, p_n) &= \ell_n\tau(y_n^\lambda(p_n - k) - y_n q) - (1 - e^{-\beta\tau})(\frac{1}{\alpha\tau}\left((\ell_n^2 y_n^{2\lambda}\alpha^2\sigma^2\tau^4)/2 - \log(\frac{e^{\beta\tau}-1}{\alpha\tau})\right)) \\ &\quad - e^{-\beta\tau}(e^{-\alpha\tau g_n^* + (\alpha y_n^\lambda \ell_n\tau^2\sigma)^2/2} - 1) \\ &= \ell_n\tau(y_n^\lambda(p_n - k) - y_n q) - (1 - e^{-\beta\tau})\left((\ell_n^2 y_n^{2\lambda}\alpha\sigma^2\tau^3)/2 - \frac{1}{\alpha\tau}\log(\frac{e^{\beta\tau}-1}{\alpha\tau})\right) \\ &\quad - e^{-\beta\tau}\left(\frac{e^{\beta\tau}-1}{\tau\alpha} - 1\right). \end{aligned} \tag{B.8}$$

Where the second equality uses the characterization in (B.7).

In order to maximize $h^3(y_n, \ell_n, p_n)$, we begin by taking the first order conditions to find the stationary point y_n^*:

$$\frac{\partial h^3(y_n^*, \ell_n, p_n)}{\partial y_n} = 0 \Leftrightarrow (y^*)^{\lambda-1}\lambda\ell_n\tau(p_n - k) - (y^*)^{2\lambda-1}\lambda\ell_n^2\tau^3\sigma^2\alpha(1 - e^{-\beta\tau}) = q\ell_n\tau. \tag{B.9}$$

In order to show that y^* is indeed a maximum of $h^3(y_n, \ell_n, p_n)$, we show in Proposition B.2 that $\frac{\partial^2 h^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{\partial y_n^2} \leq 0$, for all ℓ_n and p_n, which proves (B.5), and gives an implicit characterization of the optimal production-expenditure $y_n^*(\ell_n, p_n)$.

Putting together the results shown above, we can write

$$J_n(x_n,\ell_n,p_n) = \frac{\hat{\beta}}{\tau}\big\{(x_n - (e^{-\alpha\tau x_n} - 1) + h^3(y_n^*(\ell_n,p_n),\ell_n,p_n) + h^2(\ell_n^{d*}(\ell_n,p_n),\ell_n,p_n)\big\} \tag{B.10}$$

$$= \frac{\hat{\beta}}{\tau}\big\{(x_n - (e^{-\alpha\tau x_n} - 1) + f_n(\ell_n,p_n)\big\}, \tag{B.11}$$

where $f_n(\ell_n,p_n) = h^3(y_n^*(\ell_n,p_n),\ell_n,p_n) + h^2(\ell_n^{d*}(\ell_n,p_n),\ell_n,p_n)$. Therefore, to conclude the proof of the proposition, we need only to show that $f_n(\ell_n,p_n)$ is concave and increasing in ℓ_n, and increasing in p_n. To show this, we begin by observing that

$$h^2(\ell_n^{d*}(\ell_n,p_n),\ell_n,p_n) = \max_{\ell_n^d}[\mathbb{E}_n f_{n+1}(\ell_n + \ell_n^d, p_{n+1}) - (\ell_n^d)^+ d],$$

Where, by inductive hypothesis, $h^2(\ell_n^d,\ell_n,p_n) = \mathbb{E}_n f_{n+1}(\ell_n + \ell_n^d, p_{n+1}) - (\ell_n^d)^+ d$ is jointly concave in both ℓ_n and ℓ_n^d. Thus, because partial maximization of a jointly concave function preserves concavity, $h^2(\ell_n^{d*}(\ell_n,p_n),\ell_n,p_n)$ must be concave in ℓ_n. To show that it is increasing in ℓ_n, we observe that $\max\{\ell_n,\hat{\ell}_{n+1}\}$ is increasing in ℓ_n, and $(\hat{\ell}_{n+1}-\ell_n)^+$ is decreasing in ℓ_n, this together with the inductive hypothesis gives us the results for $h^2(\ell_n^{d*}(\ell_n,p_n),\ell_n,p_n)$. That it is increasing in p_n is a consequence of the inductive hypothesis and the fact that $\mathbb{E}(p_{n+1}|\sigma(\{p_i\}_{i\leq n}))$ is increasing in p_n.

Finally, we need only to prove that $h^3(y_n^*(\ell_n,p_n),\ell_n,p_n)$ is concave increasing in ℓ_n and increasing in p_n. We proceed by considering the first and second derivative of $h^3(y_n^*(\ell_n,p_n),\ell_n,p_n)$ with respect to ℓ_n, and the first derivative with respect to p_n.

$$\frac{dh^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{d\ell_n} = \underbrace{\frac{\partial h^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{\partial y}}_{0,\text{ by definition of } y^*} \frac{dy_n^*(\ell_n,p_n)}{d\ell_n} + \frac{\partial h^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{\partial \ell_n}$$

$$= \tau((y_n^*)^\lambda(p_n-k) - y_n^* q) - (1-e^{-\beta\tau})(\ell_n (y_n^*)^{2\lambda}\alpha\sigma^2\tau^3)$$

$$= \frac{y_n^*}{\lambda\ell_n}\left[\tau\ell_n((y_n^*)^{\lambda-1}\lambda(p_n-k) - \lambda q) - (1-e^{-\beta\tau})\lambda(\ell_n^2(y_n^*)^{2\lambda-1}\alpha\sigma^2\tau^3)\right]$$

$$= \frac{y_n^*}{\lambda\ell_n}\left[\tau\ell_n(y_n^*)^{\lambda-1}\lambda(p_n-k) - (1-e^{-\beta\tau})\lambda(\ell_n^2(y_n^*)^{2\lambda-1}\alpha\sigma^2\tau^3) - \lambda q\ell_n\tau\right]$$

$$= \frac{y_n^* q\ell_n\tau(1-\lambda)}{\lambda\ell_n}$$

$$= \frac{y_n^* q\tau(1-\lambda)}{\lambda} \geq 0.$$

The fourth equality above uses the implicit definition of y^* (B.9).

$$\frac{dh^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{dp_n} = \underbrace{\frac{\partial h^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{\partial y}}_{0,\text{ by definition of } y^*} \frac{dy_n^*(\ell_n,p_n)}{dp_n} + \frac{\partial h^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{\partial p_n}$$

$$= \tau\ell_n(y_n^*)^\lambda \geq 0.$$

This proves that $h^3(y_n^*(\ell_n,p_n),\ell_n,p_n)$ is indeed increasing in ℓ_n and p_n. Moreover, we see that $\frac{d^2h^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{d\ell_n^2} \leq 0$ if and only if $\frac{dy^*(\ell_n,p_n)}{d\ell_n} \leq 0$, which we prove in Proposition B.3. Therefore, $f_n(\ell_n,p_n)$ is both increasing in ℓ_n and concave in ℓ_n, which completes the proof of the inductive step. □

Proposition B.2 *Let $h^3(y_n,\ell_n,p_n)$ be as defined in equation (B.8), and $y^*(\ell_n,p_n)$ be the optimal production-expenditure level as defined by (B.9), then $\frac{\partial^2 h^3(y_n^*,\ell_n,p_n)}{\partial y^2} \leq 0$, for any $\ell_n \geq 0$ and p_n.*

Proof. First, let us compute the first derivative $\frac{\partial h^3(y_n^*,\ell_n,p_n)}{\partial y}$:

$$\frac{\partial h^3(y,\ell_n,p_n)}{\partial y} = y^{\lambda-1}\lambda\ell_n\tau(p_n-k) - y^{2\lambda-1}\lambda\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau}) - q\ell_n\tau.$$

From here, we can compute the desired second derivative:

$$\begin{aligned}\frac{\partial^2 h^3(y_n^*,\ell_n,p_n)}{\partial y^2} &= (y^*)^{\lambda-2}(\lambda-1)\lambda\ell_n\tau(p_n-k) - (2\lambda-1)(y^*)^{2\lambda-2}\lambda\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})\\ &= \frac{1}{y_n^*}[(y^*)^{\lambda-1}(\lambda-1)\lambda\ell_n\tau(p_n-k) - (2\lambda-1)(y^*)^{2\lambda-1}\lambda\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})]\\ &= \frac{1}{y_n^*}[(\lambda-1)q\ell_n\tau - (y^*)^{2\lambda-1}\lambda^2\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})] \le 0.\end{aligned}$$

Where the first equality uses the fact that $y^*=0$ is not a solution to equation (B.9), as long as $q>0$ (and if $q=0$, we consider the unique positive solution defined by $\left(\frac{(p_n-k)}{\ell_n\tau^2\alpha(1-e^{-\beta\tau})}\right)^{\frac{1}{\lambda}}$). Additionally, the second equality uses the definition of y^*, that implies that $(y^*)^{\lambda-1}(\lambda-1)\lambda\ell_n\tau(p_n-k) = (\lambda-1)[y^{2\lambda-1}\lambda\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})+q\ell_n\tau]$. And finally, the last inequality stems from the simple observation that $\lambda\le 1$. □

Proposition B.3 *Let $y_n^*(\ell_n,p_n)$ be the optimal production-expenditure level, as defined by (B.9), then $y_n^*(\ell_n,p_n)$ is decreasing in ℓ_n, i.e., $\frac{dy^*(\ell_n,p_n)}{d\ell_n}\le 0$.*

Proof.

We compute the derivative of $y_n^*(\ell_n,p_n)$ with respect to ℓ_n, by using the Implicit Function Theorem and the definition of y^* in (B.9).

$$\begin{aligned}\frac{dy^*(\ell_n,p_n)}{d\ell_n} &= -\frac{\frac{\partial^2 h^3(y^*,\ell_n,p_n)}{\partial y\partial\ell_n}}{\frac{\partial^2 h^3(y^*,\ell_n,p_n)}{\partial y^2}}\\ &= -\frac{(y^*)^{\lambda-1}\lambda\tau(p_n-k) - (y^*)^{2\lambda-1}\lambda 2\ell_n\tau^3\sigma^2\alpha(1-e^{-\beta\tau}) - q\tau}{\frac{\partial^2 h^3(y^*,\ell_n,p_n)}{\partial y^2}}\end{aligned}$$

But, by Proposition B.2, we know that $\frac{\partial^2 h^3(y^*,\ell_n,p_n)}{\partial y^2} \leq 0$, this implies that:

$$\begin{aligned} sign(\frac{dy^*(\ell_n,p_n)}{d\ell_n}) &= sign((y^*)^{\lambda-1}\lambda\tau(p_n-k) - (y^*)^{2\lambda-1}\lambda 2\ell_n\tau^3\sigma^2\alpha(1-e^{-\beta\tau}) - q\tau) \\ &= sign(\frac{1}{\ell_n}[(y^*)^{\lambda-1}\lambda\ell_n\tau(p_n-k) - 2(y^*)^{2\lambda-1}\lambda\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau}) \ - \ell_n q\tau)] \\ &= sign(-(y^*)^{2\lambda-1}\lambda\ell_n\tau^3\sigma^2\alpha(1-e^{-\beta\tau})) \qquad \text{by (B.9)} \end{aligned}$$

Which implies that $\frac{dy^*(\ell_n,p_n)}{d\ell_n} \leq 0$, proving the result. □

These propositions together finish the proof of Theorem 2.1. ■

We proceed to prove Theorem 2.2.

Proof of Theorem 2.2.

As a consequence of the characterization proven in Theorem 2.1, and $f_n(\ell_n,p_n)$ being increasing in ℓ_n and p_n for every $n \in \{1,\ldots,N+1\}$, we have that $J_n(x_n,\ell_n,p_n)$ must be increasing in ℓ_n and p_n. Moreover, due to this same characterizations, $J_n(x_n,\ell_n,p_n)$ is increasing in x_n if and only if $x_n + (1-e^{-\alpha x_n})$ is increasing in x_n, which can be seen by simple inspection. We need then only to prove that the value function is decreasing in q, k, and σ^2. We proceed to show by backwards induction in n that

$$\frac{\partial J_n(x_n,\ell_n,p_n)}{\partial k} \leq 0,$$

$$\frac{\partial J_n(x_n,\ell_n,p_n)}{\partial q} \leq 0,$$

$$\frac{\partial J_n(x_n,\ell_n,p_n)}{\partial \sigma^2} \leq 0.$$

When $n = N+1$, then $J_{N+1}(x_n,\ell_n,p_n) = \frac{1-e^{-\beta\tau}}{\beta\tau}\left(x_{N+1} - (e^{-\beta\tau x_{N+1}} - 1)\right)$, which implies that $\frac{\partial J_{N+1}(x_{N+1},\ell_{N+1},p_{N+1})}{\partial k} = \frac{\partial J_{N+1}(x_{N+1},\ell_{N+1},p_{N+1})}{\partial \alpha} = \frac{\partial J_{N+1}(x_{N+1},\ell_{N+1},p_{N+1})}{\partial \sigma^2} = 0$.

Now we proceed by assuming that the result holds for $n+1$, and proving that it must hold for n. Using the characterization shown in (B.10) in the proof of Proposition B.1, we can write the value function at time n as:

$$J_n(x_n, \ell_n, p_n) = \frac{\hat{\beta}}{\tau}\{(x_n - (e^{-\alpha\tau x_n} - 1) + h^3(y_n^*(\ell_n, p_n), \ell_n, p_n) + \max_{\ell_n^d \geq 0} h^2(\ell_n^d, \ell_n, p_n)\},$$

where

$$\begin{aligned} h^3(y_n^*(\ell_n, p_n), \ell_n, p_n) &= \ell_n\tau(y_n^*)^{\lambda}(p_n - k) - y_n^* q) - (1 - e^{-\beta\tau})(\ell_n^2(y_n^*)^{2\lambda}\alpha\sigma^2\tau^3)/2 \\ &\qquad - \frac{1}{\alpha\tau}\log(\frac{e^{\beta\tau} - 1}{\alpha\tau}) - e^{-\beta\tau}(\frac{e^{\beta\tau}}{\tau\alpha} - 1), \\ h^2(\ell_n^d, \ell_n, p_n) &= \mathbb{E}_n f_{n+1}(\ell_n + \ell_n^d, p_{n+1}) - (\ell_n^d)^+ d. \end{aligned}$$

Hence, if we wish to take the derivative of the value function with respect to the parameters k, q, and σ^2, we need only consider the derivatives of $h^3(y_n^*(\ell_n, p_n), \ell_n, p_n)$, and $\max_{\ell_n^d \geq 0} h^2(\ell_n^d, \ell_n, p_n)$. Let us begin by considering the latter:

$$\begin{aligned} \frac{\partial \max_{\ell_n^d \geq 0} h^2(\ell_n^d, \ell_n, p_n)}{\partial k} &= \frac{\partial \max_{\ell + n^d \geq 0} \mathbb{E}_n f_{n+1}(\ell_n + \ell_n^d, p_{n+1}) - (\ell_n^d)^+ d}{\partial k} \\ &= \frac{\partial \mathbb{E}_n f_{n+1}(\ell_n + \ell_n^d, p_{n+1}) - (\ell_n^d)^+ d}{\partial k}(\ell_n^{d*}, \ell_n, p_n) \qquad \text{Using Envelope Theorem.} \\ &= \mathbb{E}_n \frac{\partial f_{n+1}(\ell_n + \ell_n^d, p_{n+1})}{\partial k}(\ell_n^{d*}, \ell_n, p_n) \leq 0. \end{aligned}$$

Where the second equality is an application of the Envelope Theorem (see Milgrom and Segal 2002), and the final inequality comes from the inductive hypothesis. The same arguments prove that $\max_{\ell_n^d \geq 0} h^2(\ell_n^d, \ell_n, p_n)$ must be decreasing in q and σ^2. It only remains to be seen that the same is true for $h^3(y_n^*(\ell_n, p_n), \ell_n, p_n)$.

$$\frac{dh^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{dk} = \underbrace{\frac{\partial h^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{\partial y}}_{0,\text{ by definition of } y^*} \frac{dy_n^*(\ell_n,p_n)}{dk} + \frac{\partial h^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{\partial k},$$
$$= -\ell_n\tau(y_n^*)^\lambda \leq 0,$$

$$\frac{dh^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{dq} = \underbrace{\frac{\partial h^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{\partial y}}_{0,\text{ by definition of } y^*} \frac{dy_n^*(\ell_n,p_n)}{dq} + \frac{\partial h^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{\partial q},$$
$$= -\ell_n\tau(y_n^*) \leq 0,$$

$$\frac{dh^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{d\sigma^2} = \underbrace{\frac{\partial h^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{\partial y}}_{0,\text{ by definition of } y^*} \frac{dy_n^*(\ell_n,p_n)}{d\sigma^2} + \frac{\partial h^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{\partial \sigma^2},$$
$$= -\frac{1}{2}\ell_n^2\tau^3\alpha(1-e^{-\beta\tau})(y_n^*)^{2\lambda} \leq 0.$$

Where all the inequality above are readily apparent. This shows that the derivative of the value function with respect to k, σ^2, and q must be negative for n, and thus concludes the inductive proof. ■

Now we will prove Theorem 2.3, by using the same logic as in Proposition B.3.

Proof of Theorem 2.3. We wish to see that the optimal production-expenditure rate $y_n^*(\ell_n, p_n)$ as defined by (B.9), is decreasing in ℓ_n, q, k, α, and σ^2. First, notice that Proposition B.3 proves already the first result. Following the same reasoning, we will compute

$$\frac{dy^*(\ell_n, p_n)}{dq} = -\frac{\frac{\partial^2 h^3(y^*,\ell_n,p_n)}{\partial y \partial q}}{\frac{\partial^2 h^3(y^*,\ell_n,p_n)}{\partial y^2}} = -\frac{-\ell_n \tau}{\frac{\partial^2 h^3(y^*,\ell_n,p_n)}{\partial y^2}} = \frac{\ell_n \tau}{\frac{\partial^2 h^3(y^*,\ell_n,p_n)}{\partial y^2}} \leq 0.$$

Where, the numerator is always positive and the denominator was proven to be negative in Proposition B.2. We similarly compute the derivative of the optimal production-expenditure rate with respect to interest rate α and expected production cost k:

$$\frac{dy^*(\ell_n, p_n)}{d\alpha} = -\frac{\frac{\partial^2 h^3(y^*,\ell_n,p_n)}{\partial y \partial \alpha}}{\frac{\partial^2 h^3(y^*,\ell_n,p_n)}{\partial y^2}} = -\frac{-\ell_n^2 \tau^3 \sigma^2 \lambda (1 - e^{-\beta\tau})(y_n^*)^{2\lambda-1}}{\frac{\partial^2 h^3(y^*,\ell_n,p_n)}{\partial y^2}} = \frac{\ell_n^2 \tau^3 \sigma^2 \lambda (1 - e^{-\beta\tau})(y_n^*)^{2\lambda-1}}{\frac{\partial^2 h^3(y^*,\ell_n,p_n)}{\partial y^2}} \leq 0.$$

$$\begin{aligned}\frac{dy^*(\ell_n, p_n)}{dk} &= -\frac{\frac{\partial^2 h^3(y^*,\ell_n,p_n)}{\partial y \partial k}}{\frac{\partial^2 h^3(y^*,\ell_n,p_n)}{\partial y^2}} \\ &= -\frac{-\ell_n \tau \lambda (y_n^*)^{\lambda-1}}{\frac{\partial^2 h^3(y^*,\ell_n,p_n)}{\partial y^2}} \\ &= \frac{\ell_n^2 \tau^3 \sigma^2 \lambda (1-e^{-\beta\tau})(y_n^*)^{2\lambda-1}}{\frac{\partial^2 h^3(y^*,\ell_n,p_n)}{\partial y^2}} \leq 0.\end{aligned}$$

And, as before, we see that the numerator is always positive while the denominator is always negative. Finally, we compute the derivative with respect to σ^2:

$$\begin{aligned}\frac{dy^*(\ell_n, p_n)}{d\sigma^2} &= -\frac{\frac{\partial^2 h^3(y^*,\ell_n,p_n)}{\partial y \partial \sigma^2}}{\frac{\partial^2 h^3(y^*,\ell_n,p_n)}{\partial y^2}} \\ &= -\frac{-(y^*)^{2\lambda-1}\lambda \ell_n^2 \tau^3 \alpha (1-e^{-\beta\tau})}{\frac{\partial^2 h^3(y^*,\ell_n,p_n)}{\partial y^2}} \\ &= \frac{(y^*)^{2\lambda-1}\lambda \ell_n^2 \tau^3 \alpha (1-e^{-\beta\tau})}{\frac{\partial^2 h^3(y^*,\ell_n,p_n)}{\partial y^2}} \leq 0.\end{aligned}$$

Where, as before, the numerator is always positive and the denominator is always negative. This proves that the optimal production-expenditure rate is decreasing in ℓ_n, q, and σ^2.

■

We will now state and prove modularity results on $J_n(x_n, \ell_n, p_n)$ that will allow us to prove Theorem 2.4.

Proposition B.4 *The value function $J_n(x_n, \ell_n, p_n)$ is sub-modular in (ℓ_n, α), (ℓ_n, k), and (ℓ_n, σ^2), for every $n \in \{1, \ldots, N+1\}$.*

Proof. We will proceed to show by backwards induction in n that

$$\frac{\partial^2 J_n(x_n, \ell_n, p_n)}{\partial \ell_n \partial k} \leq 0,$$

$$\frac{\partial^2 J_n(x_n, \ell_n, p_n)}{\partial \ell_n \partial \alpha} \leq 0,$$

$$\frac{\partial^2 J_n(x_n, \ell_n, p_n)}{\partial \ell_n \partial \sigma^2} \leq 0.$$

When $n = N+1$, then $J_{N+1}(x_n, \ell_n, p_n) = \frac{1-e^{-\beta\tau}}{\beta\tau}\left(x_{N+1} - (e^{-\beta\tau x_{N+1}} - 1)\right)$, which implies that $\frac{\partial^2 J_{N+1}(x_{N+1}, \ell_{N+1}, p_{N+1})}{\partial \ell_{N+1} \partial k} = \frac{\partial^2 J_{N+1}(x_{N+1}, \ell_{N+1}, p_{N+1})}{\partial \ell_{N+1} \partial \alpha} = \frac{\partial^2 J_{N+1}(x_{N+1}, \ell_{N+1}, p_{N+1})}{\partial \ell_{N+1} \partial \sigma^2} = 0$. Now, as in the proof of Theorem 2.2, we proceed with the inductive step assuming the inductive hypothesis for $n+1$ and using the characterization of the value function shown in (B.10). Following the argument in the proof of Theorem 2.2, it suffices to show that $h^3(y_n^*(\ell_n, p_n), \ell_n, p_n)$ and $\max_{\ell_n^d \geq 0} h^2(\ell_n^d, \ell_n, p_n)$ are both sub-modular in (ℓ_n, α), (ℓ_n, k), and (ℓ_n, σ^2). For this, we consider the crossed derivatives and show that they are negative:

$$\begin{aligned}
\frac{\partial^2 \max_{\ell_n^d \geq 0} h^2(\ell_n^d, \ell_n, p_n)}{\partial \ell_n \partial k} &= \frac{\partial^2 \max_{\ell + n^d \geq 0} \mathbb{E}_n f_{n+1}(\ell_n + \ell_n^d, p_{n+1}) - (\ell_n^d)^+ d}{\partial \ell_n \partial k} \\
&= \frac{\partial^2 \mathbb{E}_n f_{n+1}(\ell_n + \ell_n^d, p_{n+1}) - (\ell_n^d)^+ d}{\partial \ell_n \partial k}(\ell_n^{d*}, \ell_n, p_n) \qquad \text{Using Envelope Theorem.} \\
&= \mathbb{E}_n \frac{\partial^2 f_{n+1}(\ell_n + \ell_n^d, p_{n+1})}{\partial \ell_n \partial k}(\ell_n^{d*}, \ell_n, p_n) \leq 0.
\end{aligned}$$

Where the second equality is an application of the Envelope Theorem and the last inequality is due to the inductive hypothesis. This same argument can be made for (ℓ_n, α), and (ℓ_n, σ^2). Thus, we need only to show that the crossed derivatives of $h^3(y_n^*(\ell_n, p_n), \ell_n, p_n)$ are negative. In the proof of Theorem 2.1 we have already shown that

$$\frac{dh^3(y_n^*(\ell_n, p_n), \ell_n, p_n)}{d\ell_n} = \frac{y_n^* q\tau(1-\lambda)}{\lambda} \geq 0.$$

This implies that $\frac{d^2 h^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{d\ell_n dk} \leq 0$ if and only if $\frac{dy_n^*(\ell_n,p_n)}{dk} \leq 0$, and equivalently for α, and σ^2. But we have already shown that this is the case in proving Theorem 2.3. Hence, we have $J_n(x_n, \ell_n, p_n)$ is sub-modular in (ℓ_n, k), (ℓ_n, α), and (ℓ_n, σ^2), proving the inductive step and the proposition. □

We now prove Theorem 2.4 using the modularity results from Proposition B.4.

Proof of Theorem 2.4. In Proposition B.1, we prove that $\ell_n^{d*} = (\hat{\ell}_{n+1} - \ell_n)^+$, where $\hat{\ell}_{n+1}$ solves $\frac{e^{-\beta\tau}\partial\mathbb{E}_n f_{n+1}(\ell,p_{n+1})}{\partial\ell} = d$. Additionally, in Proposition B.4, we showed that $J_n(x_n, \ell_n, p_n)$ is sub-modular in (ℓ_n, k), (ℓ_n, q), and (ℓ_n, σ^2), that by the characterization proven in Proposition B.1 implies that $f_n(\ell_n, p_n)$ is as well for every $n \in \{1, \ldots, N+1\}$. Therefore, a simple application on Topkis' theorem (Topkis 1998) shows that $\hat{\ell}_{n+1}$ must be decreasing in k, α, and σ^2 for $n \in \{1, \ldots, N\}$, which implies that ℓ_n^{d*} is decreasing for $n \in \{1, \ldots, N\}$. Finally, ℓ_{N+1}^d is always zero by definition, which concludes the proof that ℓ_n^{d*} is decreasing in k, α, and σ^2, for $n \in \{1, \ldots, N+1\}$.

■

In order to prove Theorem 2.5, we proceed to show that the modularity of $J_n(x_n, \ell_n, p_n)$, with respect to (ℓ_n, q) has the same threshold behavior.

Proposition B.5 *There exists positive functions $\tilde{q}_n^H(\ell_n)$ and $\tilde{q}_n^L(\ell_n)$ such that the value function $J_n(x_n, \ell_n, p_n)$ is **super-modular** in (ℓ_n, q), for $q \leq \tilde{q}_n^L(\ell_n)$, and **sub-modular** in (ℓ_n, q), for $q > \tilde{q}_n^H(\ell_n)$ for every $n \in \{1, \ldots, N+1\}$. Moreover, $\tilde{q}_n^H(\ell_n)$ and $\tilde{q}_n^H(\ell_n)$ are increasing in ℓ_n, σ^2, and α.*

Proof. As in the proof of Proposition B.4, we will proceed by backwards induction in n to show that there exists a $\tilde{q}(\ell_n)$ such that

$$\frac{\partial^2 J_n(x_n, \ell_n, p_n)}{\partial\ell_n \partial q} \geq 0, \text{ if } q \leq \tilde{q}_n^L(\ell_n),$$

and

$$\frac{\partial^2 J_n(x_n, \ell_n, p_n)}{\partial\ell_n \partial q} \leq 0, \text{ if } q \geq \tilde{q}_n^H(\ell_n).$$

Moreover, $\tilde{q}_n^H(\ell_n)$ and $\tilde{q}_n^L(\ell_n)$ are increasing in ℓ_n, σ^2, and α.

When $n = N+1$, $J_{N+1}(x_n, \ell_n, p_n) = \frac{1-e^{-\beta\tau}}{\beta\tau}\left(x_{N+1} - (e^{-\beta\tau x_{N+1}} - 1)\right)$, which implies that $\frac{\partial^2 J_{N+1}(x_{N+1}, \ell_{N+1}, p_{N+1})}{\partial \ell_n \partial q} = 0$. We proceed then to the inductive step, where we will assume the result holds for $n+1$.

Using the characterization shown in (B.10), we can see that it suffices to prove that $h^3(y_n^*(\ell_n, p_n), \ell_n, p_n)$ and $\max_{\ell_n^d \geq 0} h^2(\ell_n^d, \ell_n, p_n)$ both satisfy the desired property. In particular,

$$\begin{aligned}
\frac{\partial^2 \max_{\ell_n^d \geq 0} h^2(\ell_n^d, \ell_n, p_n)}{\partial \ell_n \partial q} &= \frac{\partial^2 \max_{\ell+n^d \geq 0} \mathbb{E}_n f_{n+1}(\ell_n + \ell_n^d, p_{n+1}) - (\ell_n^d)^+ d}{\partial \ell_n \partial q} \\
&= \frac{\partial^2 \mathbb{E}_n f_{n+1}(\ell_n + \ell_n^d, p_{n+1}) - (\ell_n^d)^+ d}{\partial \ell_n \partial q}(\ell_n^{d*}, \ell_n, p_n) \quad \text{Using Envelope Theorem.} \\
&= \mathbb{E}_n \frac{\partial^2 f_{n+1}(\ell_n + \ell_n^d, p_{n+1})}{\partial \ell_n \partial q}(\ell_n^{d*}, \ell_n, p_n).
\end{aligned}$$

Which, by the inductive hypothesis implies that $\frac{\partial^2 \max_{\ell_n^d \geq 0} h^2(\ell_n^d, \ell_n, p_n)}{\partial \ell_n \partial q} \geq 0$, if $q \leq \tilde{q}_{n+1}^L(\max\{\ell_n, \hat{\ell}_{n+1}\})$, and $\frac{\partial^2 \max_{\ell_n^d \geq 0} h^2(\ell_n^d, \ell_n, p_n)}{\partial \ell_n \partial q} \leq 0$, if $q \geq \tilde{q}_{n+1}^H(\max\{\ell_n, \hat{\ell}_{n+1}\})$.

Now we need only to show that the same happens for $\frac{d^2 h^3(y_n^*(\ell_n, p_n), \ell_n, p_n)}{d\ell_n dq}$. We have already shown in the proof of Proposition B.1 that

$$\frac{dh^3(y_n^*(\ell_n, p_n), \ell_n, p_n)}{d\ell_n} = \frac{y_n^* q\tau(1-\lambda)}{\lambda} \geq 0.$$

Thus, we can compute

$$\begin{aligned}
&\frac{dh^3(y_n^*(\ell_n, p_n), \ell_n, p_n)}{d\ell_n dq} \\
&= \frac{\tau(1-\lambda)}{\lambda}\left(\frac{d_n y^*(\ell_n, p_n)}{dq} q + y_n^*(\ell_n, p_n)\right) \\
&= \frac{\tau(1-\lambda)}{\lambda}\left(\frac{q\ell_n\tau}{(y_n^*)^{\lambda-2}(\lambda-1)\lambda\ell_n\tau(p_n - k) - (2\lambda-1)(y_n^*)^{2\lambda-2}\lambda\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})} + y_n^*\right).
\end{aligned}$$

Where the second equality uses the expression for $\frac{dy^*}{dq}$ proven in Theorem 2.3 (combined with the explicit form of $\frac{\partial^2 h^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{\partial y^2}$ shown in Proposition B.2).The first fraction is always positive (when $\lambda \leq 1$), which means we can analyze the sign of the crossed derivative above by looking at the sign of the expression in the parenthesis.

$$\frac{q\ell_n\tau}{(y_n^*)^{\lambda-2}(\lambda-1)\lambda\ell_n\tau(p_n-k)-(2\lambda-1)(y_n^*)^{2\lambda-2}\lambda\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})}+y_n^*=$$
$$=\frac{q\ell_n\tau+y_n^*((y_n^*)^{\lambda-2}(\lambda-1)\lambda\ell_n\tau(p_n-k)-(2\lambda-1)(y_n^*)^{2\lambda-2}\lambda\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau}))}{(y_n^*)^{\lambda-2}(\lambda-1)\lambda\ell_n\tau(p_n-k)-(2\lambda-1)(y_n^*)^{2\lambda-2}\lambda\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})}$$

As proven in Proposition B.2, the denominator will always be negative. This implies that

$$\begin{aligned}
&sign(\frac{dh^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{d\ell_n dq})\\
&=-sign(q\ell_n\tau+y_n^*((y_n^*)^{\lambda-2}(\lambda-1)\lambda\ell_n\tau(p_n-k)-(2\lambda-1)(y_n^*)^{2\lambda-2}\lambda\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})))\\
&=-sign(q\ell_n\tau+((y_n^*)^{\lambda-1}(\lambda-1)\lambda\ell_n\tau(p_n-k)-(2\lambda-1)(y_n^*)^{2\lambda-1}\lambda\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})))\\
&=-sign(q\ell_n\tau+((\lambda-1)q\ell_n\tau-\lambda^2(y_n^*)^{2\lambda-1}\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau}))) \qquad \text{by (B.5)}\\
&=-sign(q\ell_n\tau\lambda-\lambda^2(y_n^*)^{2\lambda-1}\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})))\\
&=-sign(q-\lambda(y_n^*)^{2\lambda-1}\ell_n\tau^2\sigma^2\alpha(1-e^{-\beta\tau}))\\
&=sign(\lambda(y_n^*)^{2\lambda-1}\ell_n\tau^2\sigma^2\alpha(1-e^{-\beta\tau})-q)
\end{aligned}$$

Consider then the function $s(q)=\lambda(y^*)^{2\lambda-1}\ell_n\tau^2\sigma^2\alpha(1-e^{-\beta\tau})-q$. We show that this function has exactly one zero, and that it takes positive values for q lower than this zero and negative values for higher qs. On one hand, $s(0)\geq 0$, because $y^*\geq 0$. In fact, it is easy to see that at $q=0$, (B.9) the only non-zero solution is $y^*=\left(\frac{(p_n-k)}{\ell_n\tau^2\alpha(1-e^{-\beta\tau})}\right)^{\frac{1}{\lambda}}$. On the other hand, we can see that $\lim_{q\to\infty} s(q) = -\infty$. To show this, consider

$$\begin{aligned} s(q) &= \lambda(y_n^*)^{2\lambda-1}\ell_n\tau^2\sigma^2\alpha(1-e^{-\beta\tau}) - q \\ &= \frac{(y_n^*)^{\lambda-1}}{\ell_n\tau}\Big(\lambda(y_n^*)^{\lambda}\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau}) - q\ell_n\tau(y^*)^{1-\lambda}\Big) \\ &= \frac{(y_n^*)^{\lambda-1}}{\ell_n\tau}\Big(\lambda(y_n^*)^{\lambda}\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau}) - \Big(\lambda\ell_n\tau(p_n-k) - \lambda(y_n^*)^{\lambda}\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})\Big)\Big) \quad \text{by (B.9)} \\ &= \underbrace{\frac{(y_n^*)^{\lambda-1}}{\ell_n\tau}}_{\lim_{q\to\infty}(y_n^*)^{\lambda-1}=\infty} \underbrace{\Big(2\lambda(y_n^*)^{\lambda}\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau}) - \lambda\ell_n\tau(p_n-k)\Big)}_{\lim_{q\to\infty}(y_n^*)^{\lambda}=0} \end{aligned}$$

Hence, $s(q)$ can be written as the product of two expressions, one converges to infinity and the other one to $-\lambda\ell_n\tau(p_n-k) \le 0$. We can see that $\lim_{q\to\infty}(y_n^*)^{\lambda-1} = \infty$ by taking the limit as q grows to infinity of equation (B.9):

$$\infty = \lim_{q\to\infty} q\ell_n\tau = \lim_{q\to\infty}(y_n^*)^{\lambda-1}\Big(\lambda\ell_n\tau(p_n-k) - (y^{\lambda})\lambda\ell_n^2\tau^3\sigma^2(1-e^{-\beta\tau})\Big)$$

Therefore, $\lim_{q\to\infty} s(q) = -\infty$. This implies that there must exist at least one point $\hat{q}_n^z$ that satisfies

$$s(\hat{q}_n^z) = \lambda(y_n^*(\hat{q}_n^z))^{2\lambda-1}\ell_n\tau^2\sigma^2\alpha(1-e^{-\beta\tau}) - \hat{q}_n^z = 0. \tag{B.12}$$

We show this point must be unique, by showing that at every such point $s'(\hat{q}_n^z) \le 0$.

$$\begin{aligned} \frac{ds(\hat{q}_n^z)}{dq} &= \lambda(2\lambda-1)(y_n^*(\hat{q}_n^z))^{2\lambda-2}\frac{dy_n^*(\hat{q}_n^z)(\ell_n,p_n)}{dq}\ell_n\tau^2\sigma^2\alpha(1-e^{-\beta\tau}) - 1 \\ &= \frac{\lambda(2\lambda-1)(y_n^*(\hat{q}_n^z))^{2\lambda-2}\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})}{(y_n^*(\hat{q}_n^z))^{\lambda-2}(\lambda-1)\lambda\ell_n\tau(p_n-k) - (2\lambda-1)(y_n^*(\hat{q}_n^z))^{2\lambda-2}\lambda\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})} - 1 \\ &= \frac{(y_n^*(\hat{q}_n^z))^{\lambda-2}(1-\lambda)\lambda\ell_n\tau(p_n-k) + 2(2\lambda-1)\lambda(y_n^*(\hat{q}_n^z))^{2\lambda-2}\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})}{(y_n^*(\hat{q}_n^z))^{\lambda-2}(\lambda-1)\lambda\ell_n\tau(p_n-k) - (2\lambda-1)(y_n^*(\hat{q}_n^z))^{2\lambda-2}\lambda\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})} \end{aligned}$$

We know by Proposition B.2 that the denominator will always be negative, hence, we have that:

$$
\begin{aligned}
&sign\Big(\frac{ds(\hat{q}_n^z)}{dq}\Big) \\
&= -sign\Big((y_n^*(\hat{q}_n^z))^{\lambda-2}(1-\lambda)\lambda\ell_n\tau(p_n-k) + 2(2\lambda-1)\lambda(y_n^*(\hat{q}_n^z))^{2\lambda-2}\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})\Big) \\
&= -sign\Big((y_n^*(\hat{q}_n^z))^{-1}[(y_n^*(\hat{q}_n^z))^{\lambda-1}(1-\lambda)\lambda\ell_n\tau(p_n-k) \\
&\qquad\qquad + 2(2\lambda-1)\lambda(y_n^*(\hat{q}_n^z))^{2\lambda-1}\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})]\Big) \\
&= -sign\Big((y_n^*(\hat{q}_n^z))^{\lambda-1}(1-\lambda)\lambda\ell_n\tau(p_n-k) + 2(2\lambda-1)\lambda(y_n^*(\hat{q}_n^z))^{2\lambda-1}\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})\Big) \\
&= -sign\Big((1-\lambda)[\hat{q}_n^z\ell_n\tau + \lambda(y_n^*(\hat{q}_n^z))^{2\lambda-1}\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})] \\
&\qquad + 2(2\lambda-1)\lambda(y_n^*(\hat{q}_n^z))^{2\lambda-1}\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})\Big) \\
&= -sign\Big((1-\lambda)\hat{q}_n^z\ell_n\tau + (3\lambda-1)\lambda(y_n^*(\hat{q}_n^z))^{2\lambda-1}\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})\Big) \\
&= -sign\Big((1-\lambda)\lambda(y_n^*(\hat{q}_n^z))^{2\lambda-1}\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau}) + (3\lambda-1)\lambda(y_n^*(\hat{q}_n^z))^{2\lambda-1}\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})\Big) \text{by (B.12)} \\
&= -sign\Big(2\lambda^2(y_n^*(\hat{q}_n^z))^{2\lambda-1}\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})\Big)
\end{aligned}
$$

And because this last expression is always positive, we have that $s'(\hat{q}_n^z) \leq 0$, that implies that there can only be one such zero.

We have shown that $\frac{dh^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{d\ell_n dq}$ is positive when $q \leq \hat{q}_n^z$, and negative when $q \geq \hat{q}_n^z$. Therefore, taking

$$\tilde{q}_n^H(\ell_n) = \max\{\tilde{q}_{n+1}^H(\max\{\ell_n, \hat{\ell}_{n+1}\}), \hat{q}_n^z\}, \tag{B.13}$$

$$\tilde{q}_n^L(\ell_n) = \min\{\tilde{q}_{n+1}^L(\max\{\ell_n, \hat{\ell}_{n+1}\}), \hat{q}_n^z\}, \tag{B.14}$$

satisfies the conditions, because if $q \leq \tilde{q}_n^L(\ell_n)$, both $\frac{dh^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{d\ell_n dq}$ and $\frac{\partial^2 \max_{\ell_n^d \geq 0} h^2(\ell_n^d,\ell_n,p_n)}{\partial\ell_n \partial q}$ are positive, and if $q \geq \tilde{q}_n^H(\ell_n)$, they are both negative.

Finally, because we assumed by the induction hypothesis that $tildeq_{n+1}^H(\ell_n)$ and $\tilde{q}_{n+1}^L(\ell_n)$ are increasing in ℓ_n, α, and σ^2, we need only to check that $\hat{q}_n^z$ is increasing in these three parameters to complete the induction. Thus, we take the derivatives of this expression using the implicit function theorem:

$$\frac{d\hat{q}_n^z}{d\ell_n} = -\frac{\frac{ds(\hat{q}_n^z)}{d\ell_n}}{\frac{ds(\hat{q}_n^z)}{dq}}.$$

We showed above that $\frac{ds(\hat{q}_n^z)}{dq} \leq 0$, which means that $sign(\frac{d\hat{q}_n^z}{d\ell_n}) = sign(\frac{ds(\hat{q}_n^z)}{d\ell_n})$, but

$$\begin{aligned}\frac{ds(q)}{d\ell_n} &= \tau^2\sigma^2\alpha(1-e^{-\beta\tau})\frac{d(y^*)^{2\lambda-1}\ell_n}{d\ell_n} \\ &= \tau^2\sigma^2\alpha(1-e^{-\beta\tau})\left(\ell_n(2\lambda-1)(y^*)^{2\lambda-2}\frac{dy^*}{d\ell_n} + (y^*)^{2\lambda-1}\right).\end{aligned}$$

Which implies that

$$\begin{aligned}& sign(\frac{d\hat{q}_n^z}{d\ell_n}) \\ &= sign(\frac{ds(q)}{d\ell_n}) \\ &= sign(\ell_n(2\lambda-1)(y^*)^{2\lambda-2}\frac{dy^*}{d\ell_n} + (y^*)^{2\lambda-1}) \\ &= sign(\ell_n(2\lambda-1)\frac{dy^*}{d\ell_n} + y^*) && ((y^*)^{2\lambda-2} \geq 0) \\ &= sign(\frac{\ell_n(2\lambda-1)(y^*)^{2\lambda-1}\lambda\ell_n\tau^3\sigma^2\alpha(1-e^{-\beta\tau})}{[(y^*)^{\lambda-2}(\lambda-1)\lambda\ell_n\tau(p_n-k) - (2\lambda-1)(y^*)^{2\lambda-2}\lambda\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})]} + y^*) && \text{by Prop. B.3} \\ &= sign(\frac{(y^*)^{\lambda-1}(\lambda-1)\lambda\ell_n\tau(p_n-k)}{[(y^*)^{\lambda-2}(\lambda-1)\lambda\ell_n\tau(p_n-k) - (2\lambda-1)(y^*)^{2\lambda-2}\lambda\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})]}).\end{aligned}$$

The final expression is always positive, because the denominator is exactly $\frac{\partial^2 h^3(y_n^*(\ell_n,p_n),\ell_n,p_n)}{\partial y^2}$, which we proved in Proposition B.2 to be negative, and the numerator is always negative when $\lambda \leq 1$.

Similarly, we can compute:

$$
\begin{aligned}
& sign(\frac{d\hat{q}_n^z}{d\ell_n}) \\
& = sign(\frac{ds(q)}{d\alpha}) \\
& = sign(\alpha(2\lambda-1)(y^*)^{2\lambda-2}\frac{dy^*}{d\alpha} + (y^*)^{2\lambda-1}) \\
& = sign(\alpha(2\lambda-1)\frac{dy^*}{d\alpha} + y^*) && ((y^*)^{2\lambda-2} \geq 0) \\
& = sign(\frac{\alpha(2\lambda-1)(y^*)^{2\lambda-1}\lambda\ell_n^2\tau^3\sigma^2(1-e^{-\beta\tau})}{[(y^*)^{\lambda-2}(\lambda-1)\lambda\ell_n\tau(p_n-k) - (2\lambda-1)(y^*)^{2\lambda-2}\lambda\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})]} + y^*) && \text{by Theo. 2.3} \\
& = sign(\frac{(y^*)^{\lambda-1}(\lambda-1)\lambda\ell_n\tau(p_n-k)}{[(y^*)^{\lambda-2}(\lambda-1)\lambda\ell_n\tau(p_n-k) - (2\lambda-1)(y^*)^{2\lambda-2}\lambda\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})]}),
\end{aligned}
$$

and

$$
\begin{aligned}
& sign(\frac{d\hat{q}_n^z}{d\ell_n}) \\
& = sign(\frac{ds(q)}{d\sigma^2}) \\
& = sign(\sigma^2(2\lambda-1)(y^*)^{2\lambda-2}\frac{dy^*}{d\sigma^2} + (y^*)^{2\lambda-1}) \\
& = sign(\sigma^2(2\lambda-1)\frac{dy^*}{d\sigma^2} + y^*) && ((y^*)^{2\lambda-2} \geq 0) \\
& = sign(\frac{\sigma^2(2\lambda-1)(y^*)^{2\lambda-1}\lambda\ell_n^2\tau^3\alpha(1-e^{-\beta\tau})}{[(y^*)^{\lambda-2}(\lambda-1)\lambda\ell_n\tau(p_n-k) - (2\lambda-1)(y^*)^{2\lambda-2}\lambda\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})]} + y^*) && \text{by Theo. 2.3} \\
& = sign(\frac{(y^*)^{\lambda-1}(\lambda-1)\lambda\ell_n\tau(p_n-k)}{[(y^*)^{\lambda-2}(\lambda-1)\lambda\ell_n\tau(p_n-k) - (2\lambda-1)(y^*)^{2\lambda-2}\lambda\ell_n^2\tau^3\sigma^2\alpha(1-e^{-\beta\tau})]}).
\end{aligned}
$$

Where in both cases we obtain the same expression as before, which is always positive. This proves that the thresholds are always increasing in ℓ_n, α, and σ^2, and concludes the inductive proof of the proposition. □

Proof of Theorem 2.5 Using the same arguments as in the proof of Theorem 2.4, and the modularity results proven in Proposition B.5, we obtain the proof of this Theorem. ■

Bibliography

Abman R, Carney C (2020) Agricultural productivity and deforestation: Evidence from input subsidies and ethnic favoritism in malawi. *Journal of Environmental Economics and Management* 103:102342.

Akkaya D, Bimpikis K, Lee H (2021) Government interventions to promote agricultural innovation. *Manufacturing & Service Operations Management* 23(2):437–452, URL `http://dx.doi.org/10.1287/msom.2019.0834`.

Alcott B (2005) Jevons' paradox. *Ecological economics* 54(1):9–21.

Alizamir S, Iravani F, Mamani H (2019) An analysis of price vs. revenue protection: Government subsidies in the agriculture industry. *Management Science* 65(1):32–49, URL `http://dx.doi.org/10.1287/mnsc.2017.2927`.

Angelsen A, Kaimowitz D (1999) Rethinking the causes of deforestation: lessons from economic models. *The world bank research observer* 14(1):73–98.

Angelsen A, Kaimowitz D (2001) *Agricultural technologies and tropical deforestation* (CABi).

Angelsen A, van Soest D, Kaimowitz D, Bulte E (2001) Technological change and deforestation: a theoretical overview. *Agricultural technologies and tropical deforestation* 19–34.

Aravindakshan S, Krupnik TJ, Amjath-Babu T, Speelman S, Tur-Cardona J, Tittonell P, Groot JC (2021) Quantifying farmers' preferences for cropping systems intensification: A choice experiment approach applied in coastal bangladesh's risk prone farming systems. *Agricultural Systems* 189:103069, ISSN 0308-521X, URL `http://dx.doi.org/https://doi.org/10.1016/j.agsy.2021.103069`.

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

Asfaw S, Shiferaw B, Simtowe F, Lipper L (2012) Impact of modern agricultural technologies on smallholder welfare: Evidence from tanzania and ethiopia. *Food Policy* 37(3):283–295, ISSN 0306-9192, URL `http://dx.doi.org/https://doi.org/10.1016/j.foodpol.2012.02.013`.

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.

Bertsimas D, Farias VF, Trichakis N (2013) Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research* 61(1):73–87, URL `http://dx.doi.org/10.1287/opre.1120.1138`.

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

Bouchery Y, Corbett CJ, Fransoo JC, Tan T (2016) *Sustainable supply chains: A research-based textbook on operations and strategy*, volume 4 (Springer).

Boyabatlı O, Nasiry J, Zhou YH (2019) Crop planning in sustainable agriculture: Dynamic farmland allocation in the presence of crop rotation benefits. *Management Science* 65(5):2060–2076, URL `http://dx.doi.org/10.1287/mnsc.2018.3044`.

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

Bruun TB, de Neergaard A, Burup ML, Hepp CM, Larsen MN, Abel C, Aumtong S, Magid J, Mertz O (2017) Intensification of upland agriculture in thailand: Development or degradation? *Land*

Degradation & Development 28(1):83–94, URL http://dx.doi.org/https://doi.org/10.1002/ldr.2596.

Bélanger G, Walsh JR, Richards JE, Milburn PH, Ziadi N (2000) Comparison of three statistical models describing potato yield response to nitrogen fertilizer. *Agronomy Journal* 92(5):902–908, URL http://dx.doi.org/https://doi.org/10.2134/agronj2000.925902x.

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

Carlson K, Heilmayr R, Gibbs H, Noojipady P, Burns D, Morton D, Walker N, Paoli G, Kremen C (2018) Effect of oil palm sustainability certification on deforestation and fire in indonesia. *Proceedings of the National Academy of Sciences* 115(1):121–126.

Cerrato ME, Blackmer AM (1990) Comparison of models for describing; corn yield response to nitrogen fertilizer. *Agronomy Journal* 82(1):138–143, URL http://dx.doi.org/https://doi.org/10.2134/agronj1990.00021962008200010030x.

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

Chintapalli P, Tang CS (2018) The impact of crop minimum support prices on crop production and farmer welfare. *Available at SSRN 3262407* .

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.

Collins D, Rutherford S, Morduch J (2009) *Portfolios of the Poor* (Princeton University Press).

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

Dawande M, Gavirneni S, Mehrotra M, Mookerjee V (2013) Efficient distribution of water between head-reach and tail-end farms in developing countries. *Manufacturing & Service Operations Management* 15(2):221–238, URL http://dx.doi.org/10.1287/msom.1120.0414.

de Zegher JF, Iancu DA, Plambeck EL (2018) Sustaining smallholders and rainforests by eliminating payment delay in a commodity supply chain—it takes a village Submitted to Management Science.

Djurfeldt G, Holmén H, Jirström M, Larsson R (2005) *The African food crisis: lessons from the Asian Green Revolution.* (Wallingford, UK: CABI Publishing).

DoorDash (2019) DoorDash Challenges. `https://help.doordash.com/dashers/s/challenges?language=en_US`, [Web page accessed 14-Nov-2019].

Duflo E, Banerjee A (2011) *Poor Economics* (New York City, NY: PublicAffairs).

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.

Federgruen A, Lall U, Şimşek AS (2019) Supply chain analysis of contract farming. *Manufacturing & Service Operations Management* 21(2):361–378, URL `http://dx.doi.org/10.1287/msom.2018.0735`.

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].

Garrett RD, Lambin EF, Naylor RL (2013) Land institutions and supply chain configurations as determinants of soybean planted area and yields in brazil. *Land Use Policy* 31:385–396, ISSN 0264-8377, URL `http://dx.doi.org/https://doi.org/10.1016/j.landusepol.2012.08.002`, themed Issue 1-Guest Editor Romy GreinerThemed Issue 2- Guest Editor Davide Viaggi.

Gatto M, Wollni M, Asnawi R, Qaim M (2017) Oil palm boom, contract farming, and rural economic development: village-level evidence from indonesia. *World Development* 95:127–140.

Geist HJ, Lambin EF (2002) Proximate causes and underlying driving forces of tropical deforestationtropical forests are disappearing as the result of many pressures, both local and regional, acting in various combinations in different geographical locations. *BioScience* 52(2):143–150.

Ghosh P, Mookherjee D, Ray D (2000) Credit rationing in developing countries: An overview of the theory. Mookherjee D, Ray D, eds., *A Reader in Development Economics* (London: Blackwell).

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).

GrubHub (2019) Guaranteed Hourly Minimums. `https://driver-support.grubhub.com/hc/en-us/sections/360004883032-Guaranteed-Hourly-Minimums`, [Web page accessed 14-Nov-2019].

Hagin J (1960) On the shape of the yield curve. *Plant and Soil* 12(3):285–296, ISSN 0032079X, 15735036, URL `http://www.jstor.org/stable/42931846`.

Harvey CA, Rakotobe ZL, Rao NS, Dave R, Razafimahatratra H, Rabarijohn RH, Rajaofara H, MacKinnon JL (2014) Extreme vulnerability of smallholder farmers to agricultural risks and climate change in madagascar. *Philosophical Transactions of the Royal Society B: Biological Sciences* 369(1639):20130089, URL `http://dx.doi.org/10.1098/rstb.2013.0089`.

Hazell PB (2009) *The Asian green revolution*, volume 911 (Intl Food Policy Res Inst).

Houghton R (2012) Carbon emissions and the drivers of deforestation and forest degradation in the tropics. *Current Opinion in Environmental Sustainability* 4(6):597–603.

Hu M, Liu Y, Wang W (2019) Socially beneficial rationality: The value of strategic farmers, social entrepreneurs, and for-profit firms in crop planting decisions. *Management Science* 65(8):3654–3672, URL `http://dx.doi.org/10.1287/mnsc.2018.3133`.

Iancu DA, Trichakis N (2014) Fairness and efficiency in multiportfolio optimization. *Operations Research* 62(6):1285–1301, URL `http://dx.doi.org/10.1287/opre.2014.1310`.

IFAD U (2013) Smallholders, food security and the environment. *Rome: International Fund for Agricultural Development* 29.

India-Agro-Net (2021a) Crop cultivation guide: Ginger. Online, accessed May 18th, 2021. `https://www.indiaagronet.com/horticulture/CONTENTS/ginger.htm`.

India-Agro-Net (2021b) Crop cultivation guide: Maize. Online, accessed May 18th, 2021. `https://www.indiaagronet.com/indiaagronet/crop%20info/maize.htm`.

Jayasuriya S, et al. (2001) Agriculture and deforestation in tropical asia: an analytical framework. *Agricultural technologies and tropical deforestation* 1:317–335.

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

Joffre OM, Poortvliet PM, Klerkx L (2018) Are shrimp farmers actual gamblers? an analysis of risk perception and risk management behaviors among shrimp farmers in the mekong delta. *Aquaculture* 495:528–537, ISSN 0044-8486, URL http://dx.doi.org/https://doi.org/10.1016/j.aquaculture.2018.06.012.

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.

Kalkanci B, Rahmani M, Toktay LB (2019) The role of inclusive innovation in promoting social sustainability. *Production and Operations Management* 28(12):2960–2982.

Ketterings QM, Tri Wibowo T, van Noordwijk M, Penot E (1999) Farmers' perspectives on slash-and-burn as a land clearing method for small-scale rubber producers in sepunggur, jambi province, sumatra, indonesia. *Forest Ecology and Management* 120(1):157–169, ISSN 0378-1127, URL http://dx.doi.org/https://doi.org/10.1016/S0378-1127(98)00532-5.

Kissinger G, Herold M, De Sy V (2012) Drivers of deforestation and forest degradation: a synthesis report for redd+ policymakers. Technical report, Citeseer.

Kleindorfer PR, Singhal K, Van Wassenhove LN (2005) Sustainable operations management. *Production and Operations Management* 14(4):482–492, URL http://dx.doi.org/https://doi.org/10.1111/j.1937-5956.2005.tb00235.x.

Koutsoupias E, Papadimitriou C (2009) Worst-case equilibria. *Computer Science Review* 3(2):65 – 69.

Kroeger A, Bakhtary H, Haupt F, Streck C (2017) *Eliminating Deforestation from the Cocoa Supply Chain* (World Bank), URL http://dx.doi.org/10.1596/26549.

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.

Lee HL, Tang CS (2018) Socially and environmentally responsible value chain innovations: New operations management research opportunities. *Management Science* 64(3):983–996, URL http://dx.doi.org/10.1287/mnsc.2016.2682.

Levi R, Singhvi S, Zheng Y (2020) Economically motivated adulteration in farming supply chains. *Management Science* 66(1):209–226, URL http://dx.doi.org/10.1287/mnsc.2018.3215.

Li L, Shubik M, Sobel M (2013) Control of dividends, capital subscriptions, and physical inventories. *Management Science* 59(5):1107–1124.

Li W, Ashlagi I, Lo I (2020) Simple and approximately optimal contracts for payment for ecosystem services. *Available at SSRN 3754041* .

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].

Maertens M, Zeller M, Birner R (2006) Sustainable agricultural intensification in forest frontier areas. *Agricultural Economics* 34(2):197–206, URL http://dx.doi.org/https://doi.org/10.1111/j.1574-0864.2006.00118.x.

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

Mason CF, Plantinga AJ (2013) The additionality problem with offsets: Optimal contracts for carbon sequestration in forests. *Journal of Environmental Economics and Management* 66(1):1–14, ISSN 0095-0696, URL http://dx.doi.org/https://doi.org/10.1016/j.jeem.2013.02.003.

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

Milgrom P, Segal I (2002) Envelope theorems for arbitrary choice sets. *Econometrica* 70(2):583–601.

Nelson GC, Valin H, Sands RD, Havlík P, Ahammad H, Deryng D, Elliott J, Fujimori S, Hasegawa T, Heyhoe E, Kyle P, Von Lampe M, Lotze-Campen H, Mason d'Croz D, van Meijl H, van der Mensbrugghe D, Müller C, Popp A, Robertson R, Robinson S, Schmid E, Schmitz C, Tabeau A, Willenbockel D (2014) Climate change effects on agriculture: Economic responses to biophysical shocks. *Proceedings of the National Academy of Sciences* 111(9):3274–3279, ISSN 0027-8424, URL http://dx.doi.org/10.1073/pnas.1222465110.

New York City Department of Transportation (2018) NYC Mobility Report 2018. http://www.nyc.gov/html/dot/downloads/pdf/mobility-report-2018-print.pdf.

Ning J, Sobel M (2017) Production and capacity management with internal financing. *Manufacturing & Service Operations Management* .

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].

Pelletier J, Ngoma H, Mason NM, Barrett CB (2020) Does smallholder maize intensification reduce deforestation? evidence from zambia. *Global Environmental Change* 63:102127, ISSN 0959-3780, URL http://dx.doi.org/https://doi.org/10.1016/j.gloenvcha.2020.102127.

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.

Qi J (2017) Mitigating delays and unfairness in appointment systems. *Management Science* 63(2):566–583.

Rapsomanikis G (2015) The economic lives of smallholder farmers: An analysis based on household data from nine countries. *Food and Agriculture Organization of the United Nations, Rome* .

Rashid S, Dorosh PA, Malek M, Lemma S (2013) Modern input promotion in sub-saharan africa: insights from asian green revolution. *Agricultural Economics* 44(6):705–721, URL http://dx.doi.org/https://doi.org/10.1111/agec.12083.

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.

Shephard RW, Färe R (1974) The law of diminishing returns. *Zeitschrift für Nationalökonomie /*

Journal of Economics 34(1/2):69–90, ISSN 00443158, 23048360, URL http://www.jstor.org/stable/41797657.

Shively GE, Pagiola S (2004) Agricultural intensification, local labor markets, and deforestation in the philippines. *Environment and Development Economics* 241–266.

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].

Tilman D, Balzer C, Hill J, Befort BL (2011) Global food demand and the sustainable intensification of agriculture. *Proceedings of the National Academy of Sciences* 108(50):20260–20264, ISSN 0027-8424, URL http://dx.doi.org/10.1073/pnas.1116437108.

Topkis D (1998) *Supermodularity and Complementarity* (Princeton University Press).

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].

Verheye W (2010) Growth and production of oil palm. Verheye WH, ed., *Land use, land cover and soil sciences*, 24, Encyclopedia of Life Support Systems (EOLSS) (UNESCO-EOLSS Publishers).

Villoria NB, Golub A, Byerlee D, Stevenson J (2013) Will yield improvements on the forest frontier reduce greenhouse gas emissions? a global analysis of oil palm. *American Journal of Agricultural Economics* 95(5):1301–1308, URL http://dx.doi.org/https://doi.org/10.1093/

ajae/aat034.

Wolgin JM (1975) Resource allocation and risk: A case study of smallholder agriculture in kenya. *American Journal of Agricultural Economics* 57(4):622–630, URL http://dx.doi.org/https://doi.org/10.2307/1238880.

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.