

**THREE ESSAYS IN THE INTERFACE OF
OPTIMIZATION WITH MECHANISM DESIGN,
NONEXCLUSIVE COMPETITION, AND PROPHET
INEQUALITIES**

A DISSERTATION SUBMITTED TO
THE GRADUATE SCHOOL OF ENGINEERING AND SCIENCE
OF BILKENT UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR
THE DEGREE OF
DOCTOR OF PHILOSOPHY
IN
INDUSTRIAL ENGINEERING

By
Halil İbrahim Bayrak
September 2022

Three Essays in the Interface of Optimization with Mechanism Design,
Nonexclusive Competition, and Prophet Inequalities
By Halil İbrahim Bayrak
September 2022

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

Mustafa Çelebi Pınar(Advisor)

Nuh Aygün Dalkıran(Co-Advisor)

Savaş Dayanık

Tarık Kara

Ethem Akyol

İsmail Sağlam

Approved for the Graduate School of Engineering and Science:

Orhan Arıkan
Director of the Graduate School

ABSTRACT

THREE ESSAYS IN THE INTERFACE OF OPTIMIZATION WITH MECHANISM DESIGN, NONEXCLUSIVE COMPETITION, AND PROPHET INEQUALITIES

Halil İbrahim Bayrak

Ph.D. in Industrial Engineering

Advisor: Mustafa Çelebi Pınar

Co-Advisor: Nuh Aygün Dalkıran

September 2022

Mechanism Design. We consider the mechanism design problem of a principal allocating a single good to one of several agents without monetary transfers. Each agent desires the good and uses it to create value for the principal. We designate this value as the agent's private type. Even though the principal does not know the agents' types, she can verify them at a cost. The allocation of the good thus depends on the agents' self-declared types and the results of any verification performed, and the principal's payoff matches her allocation value minus the verification costs. It is known that when the agents' types are independent, a favored-agent mechanism maximizes her expected payoff. However, this result relies on the unrealistic assumptions that the agents' types follow known independent probability distributions. We assume that the agents' types are governed by an ambiguous joint probability distribution belonging to a commonly known ambiguity set and that the principal maximizes her worst-case expected payoff. We consider three types of ambiguity sets: (i) support-only ambiguity sets, which contain all distributions supported on a rectangle, (ii) Markov ambiguity sets, characterized through first-order moment bounds, and (iii) Markov with independence ambiguity sets. For each of these ambiguity sets, we show that a favored-agent mechanism, which we characterize implicitly, is optimal and also Pareto-robustly optimal. The optimal choices of the favored agent and the threshold do not depend on the verification costs in all three cases.

Nonexclusive Competition. A freelancer with a time constraint faces offers from multiple identical parties. The quality of the service provided by the freelancer can be high or low

and is only known by the freelancer. The freelancer's time cost is strictly increasing and convex. We show that a pure-strategy equilibrium exists if and only if the preferences of the high-type freelancer satisfy one of the following two distinct conditions: (i) the high-type freelancer does not prefer providing his services for a price equal to the expected quality at the no-trade point; (ii) the high-type freelancer prefers providing his services for a price equal to the expected quality at any feasible trade point. If (i) holds, then in equilibrium, the high-type freelancer does not trade, whereas the low-type may not trade, trade efficiently, or exhaust all of his capacity. Moreover, the buyers make zero profit from each of their traded contracts. If (ii) holds, then both types of the freelancer trade at the capacity in equilibrium. Furthermore, the buyers make zero expected profit with cross-subsidization. In any equilibrium, the aggregate equilibrium trades are unique.

Prophet Inequalities. Prophet inequalities bound the expected reward obtained in a class of stopping problems by the optimal reward of the corresponding offline problem. We show how to obtain prophet inequalities for a large class of stopping problems associated with selecting a point in a polyhedron. Our approach utilizes linear programming tools and is based on a reduced form representation of the stopping problem. We illustrate the usefulness of our approach by re-establishing three different prophet inequality results from the literature. (i) For polymatroids with nonnegative coefficients in their unique Minkowski sum of simplices, we prove the $\frac{1}{2}$ -prophet inequality. (ii) We prove the $\frac{1}{n}$ -prophet inequality when there are n stages, the stages have dependently distributed rewards, and we are restricted to choosing a strategy from an arbitrary polyhedron. (iii) When the feasible set of strategies can be described via K different constraints, we obtain the $\frac{1}{K+1}$ -prophet inequality.

Keywords: Mechanism design, Costly verification, Distributionally robust optimization, Ambiguity aversion & Adverse Selection, Competing Mechanisms, Nonexclusivity, Labor Markets & Stopping Problems, Prophet Inequalities.

ÖZET

MEKANİZMA TASARIMI, MÜNHASİR OLMAYAN REKABET VE KAHİN EŞİTSİZLİKLERİ İLE ENİYİLEME ARAYÜZÜNDE ÜÇ DENEME

Halil İbrahim Bayrak

Endüstri Mühendisliği, Doktora

Tez Danışmanı: Mustafa Çelebi Pınar

İkinci Tez Danışmanı: Nuh Aygün Dalkıran

Eylül 2022

Mekanizma Tasarımı. Para alış verişi olmaksızın birkaç adaydan birine tek bir mal tahsis etmek üzerine bir mekanizma tasarım problemini ele alıyoruz. Her aday mala sahip olmak ister ve bunu mekanizma tasarımcısı için değer yaratmak için kullanır. Bu değer, her adaya özel ve mahremdir. Mekanizma tasarımcısı bu değerleri bilmese de, bir ücret karşılığında onları teftiş edebilir. Dolayısıyla, malın tahsisi, adayların kendi beyan ettikleri değerlere ve gerçekleştirilen herhangi bir doğrulamanın sonuçlarına bağlıdır. Mekanizma tasarımcısının kazancı, tahsis değerinden teftiş maliyetleri çıkarılınca bulunan değere eşittir. Adayların değerleri bağımsız dağıtılan rasgele değişkenler olduğunda, tercih edilen aday mekanizmasının beklenen getiriyi enbüyüklediği bilinmektedir. Ancak bu sonuç, aday değerlerinin bilinen ve bağımsız olasılık dağılımlarına göre belirlendiğine dair gerçekçi olmayan varsayımlara dayanmaktadır. Biz ise aday değerlerinin bağlı olduğu olasılık dağılımının herkesçe bilinen bir belirsizlik kümesine ait olduğunu, ve mekanizma tasarımcısının en kötü durumda gerçekleşen beklenen getirisini enbüyüklemek istediğini varsayıyoruz. Üç tür belirsizlik kümesini ele alıyoruz: (i) bir dikdörtgen içindeki tüm dağılımları içeren belirsizlik kümeleri, (ii) bir dikdörtgen içinde olan, adayların beklenen değerlerini de sınırlayan belirsizlik kümeleri ve (iii) bir dikdörtgen içinde olan, adayların beklenen değerlerinin sınırlandığı ve bağımsız olduğu belirsizlik kümeleri. Bu belirsizlik kümelerinin her biri için, eniyi ve ayrıca Pareto-gürbüz eniyi olan bir tercih edilen aday mekanizması olduğunu gösteriyoruz. Her üç durumda da, Pareto-gürbüz eniyi mekanizmayı tanımlayan aday ve eşik değerleri teftiş maliyetlerinden bağımsız olarak seçilebilir.

Münhasır Olmayan Rekabet. Zaman kısıtlaması olan bir serbest çalışan, birden fazla

potansiyel müşteriden gelen tekliflerle karşı karşıya kalır. Serbest çalışan tarafından verilen hizmetin kalitesi yüksek veya düşük olabilir, ve bu bilgi sadece serbest çalışan tarafından bilinir. Serbest çalışanın maliyeti dışbükeydir ve çalışma süresi arttıkça kesinlikle artar. Bir saf strateji dengesinin ancak ve ancak yüksek tipte serbest çalışanın tercihleri aşağıdaki iki farklı koşuldandır birini karşılıyorsa var olduğunu gösteriyoruz: (i) yüksek tip serbest çalışan, ticaret yapılmayan noktada hizmetlerini beklenen kaliteye eşit bir fiyata sunmayı tercih etmiyor; (ii) yüksek tip serbest çalışan, herhangi bir uygun ticaret noktasında hizmetlerini beklenen kaliteye eşit bir fiyata sunmayı tercih ediyor. Eğer (i) tutarsa, o zaman dengede, yüksek tip serbest çalışan ticaret yapmazken, düşük tip ticaret yapmayabilir, verimli ticaret yapabilir veya tüm kapasitesini tüketebilir. Ayrıca, alıcılar işlem gören sözleşmelerinin her birinden sıfır kar elde eder. Eğer (ii) tutarsa, o zaman dengede, her iki serbest meslek türü de kapasitede ticaret yapar. Ayrıca, alıcılar yüksek tipte yapılan alışverişten kar ederken, düşük tipte olan alışverişten zarar ederler. Her alıcının beklenen kazancı ise sıfırdır. Herhangi bir dengede, toplam denge işlemlerinin alabileceği tek bir değer vardır.

Kahin Eşitsizlikleri. Kahin eşitsizlikleri, durdurma problemlerinde elde edilen beklenen ödül, karşılık gelen çevrimdışı problemin eniyi ödülünü kullanarak sınırlar. Bir çokyüzlüden strateji seçilmesi de dahil, geniş bir durdurma problemi sınıfı için kahin eşitsizliklerinin nasıl elde edileceğini gösteriyoruz. Doğrusal programlama araçlarını kullanan çözüm yöntemimiz, durdurma probleminin indirgenmiş temsiline dayanmaktadır. Literatürden üç farklı kahin eşitsizliği sonucunu yeniden elde ederek yaklaşımımızın yararlılığını gösteriyoruz. (i) Eşsiz Minkowski Simpleks toplamında eksi değerler olmayan polimatroid'ler için $\frac{1}{2}$ -kahin eşitsizliğini ispatlıyoruz. (ii) Evre sayısı n olduğunda, evrelerdeki ödüller birbirlerine bağımlı dağıtıldığında ve herhangi bir çokyüzlüden strateji seçme kısıtı olduğunda, $\frac{1}{n}$ -kahin eşitsizliğini ispatlıyoruz. (iii) Seçilebilen strateji kümesi K adet kısıtlama ile tanımlanabildiğinde, $\frac{1}{K+1}$ -kahin eşitsizliğini elde ediyoruz.

Anahtar sözcükler: Mekanizma tasarımı, Masraflı Teftiş, Dağıtım Açısından Gürbüz Eniyileme, Belirsizlikten Kaçınma & Ters Seçim, Rekabet Mekanizmaları, Münhasır Olmayan Rekabet, İşgücü Piyasaları & Durdurma problemleri, Kahin eşitsizlikleri.

Acknowledgement

First and foremost, I would like to express my deepest gratitude and appreciation to my advisor Prof. Mustafa Çelebi Pınar, for his constant guidance, support, and encouragement. I have learned so much from him, both scientifically and personally. I am also thankful for having Asst. Prof. Nuh Aygün Dalkıran as my co-advisor. Our discussions made this interdisciplinary thesis possible. I consider myself lucky to have such mentors. I want to thank Prof. Rakesh Vohra for giving me the opportunity to visit the University of Pennsylvania under his supervision. I am highly grateful for his guidance in Chapter 4 of this thesis.

I gratefully acknowledge that my visit to the University of Pennsylvania was financially supported by The Scientific and Technological Research Council of Turkey (TÜBİTAK) with grant number BİDEB-2214A.

I should also mention and thank Asst. Prof. Çağıl Koçyiğit and Prof. Daniel Kuhn. I had the pleasure of collaborating with them on the problem presented in Chapter 2. Closing the scientific part of this acknowledgement, I would like to thank the members of my thesis committee, Prof. Savaş Dayanık, Asst. Prof. Tarık Kara, Asst. Prof. Ethem Akyol and Prof. İsmail Sağlam for their time and valuable comments.

Through my graduate studies, I had a very supportive group of close friends who believed in me more than I did. Thank you Cansu & Cemal İlhan, Kamyar & Nihal Kargar, Halenur & İrfan Mahmutoğulları and Haşim Özlü for being there the whole time.

I am much grateful for my family. My parents, Meva and Ahmet, deserve special thanks for their love and care. Their unconditional support made me the man that I am today. I would also like to thank my sisters Sümeyye, Zeynep, and my brother Talha for bearing with me. Lastly, I am much thankful for my lovely, soon-to-be fiancée Ümmü Seleme Nizam. I only wish I had met you much before so that we could have shared much more. Thank you for loving me.

Contents

1	Introduction	1
1.1	The Contributions and the Structure of the Dissertation	2
1.1.1	Statement of Originality	4
2	Distributionally Robust Optimal Allocation with Costly Verification	5
2.1	Introduction	5
2.2	Problem Statement and Preliminaries	10
2.3	Support-Only Ambiguity Sets	17
2.4	Markov Ambiguity Sets	31
2.5	Markov Ambiguity Sets with Independent Types	48
3	Nonexclusive Competition for a Freelancer under Adverse Selection	70
3.1	Introduction	70
3.2	The Model	75
3.3	Equilibrium Characterization	77
3.4	The Main Results	87
3.5	Concluding Remarks	93
4	Prophet Inequalities & Polyhedrons	96
4.1	Introduction	96
4.2	Model	99
4.3	Results	103

4.3.1	Polymatroids	103
4.3.2	Correlated Rewards	107
4.3.3	K Many Constraints	112
4.4	Concluding Remarks	114

List of Figures

- 2.1 Partition of the type space \mathcal{T} , (Support-Only). 23
- 2.2 Partition of the type space \mathcal{T} , (Markov). 38
- 2.3 Partition of the type space \mathcal{T} , (Markov with Independence). 57

List of Tables

3.1 Comparison with Attar et al. [1, 2] 74

Chapter 1

Introduction

This thesis consists of three chapters that share little to no common ground except for using optimization and linear programming tools. Chapter 2 presents and solves a robust mechanism design problem. In general, a mechanism designer owns an item desired by several agents who privately know their valuation for the item. The aim is to find an allocation rule that elicits privately held information and performs well with respect to some chosen objective. Robustness comes into play when there is ambiguity regarding the distribution governing privately held information. Hence, it is also in the interest of the mechanism designer to find an allocation rule that performs well even when the distribution is adversely selected. We use mechanism design and robust optimization tools to deliver our results for this part. Chapter 3 focuses on labor markets where several buyers compete nonexclusively to acquire the services of a freelancer under adverse selection. Adverse selection means that the freelancer privately knows his service quality, and a high-quality service comes with a high price. On the other hand, nonexclusive competition means that the buyers cannot force the freelancer to work exclusively for them. Hence, as long as his capacity allows, the freelancer can contract with several buyers simultaneously. In this setting, we characterize the aggregate equilibrium trades by studying certain deviations of the

buyers. In Chapter 4, we show that linear programming tools can be used to drive prophet inequalities more straightforwardly. Prophet inequalities are approximation guarantees for optimal stopping problems, which compare the performance of the optimal solution to that of complete information.

1.1 The Contributions and the Structure of the Dissertation

Here, we give a brief presentation of each chapter and their contributions. We will discuss and contrast them with the existing literature in more detail in their respective chapters.

Chapter 2 is a result of collaborations with Dr. Çağıl Koçyiğit, Dr. Daniel Kuhn, and my advisor Dr. Mustafa Çelebi Pınar. We consider a principal who wants to allocate a good to one of several agents without using monetary transfers. Each agent derives strictly positive utility from owning the good and privately knows the value he generates for the principal if he is allocated with the good. The principal can verify any agent's report at a cost, which will perfectly reveal the agent's type. The principal aims to design a mechanism that maximizes her value from the allocation minus verification costs. When the agents' types are modeled as independent random variables governed by a commonly known probability distribution, favored-agent mechanisms solve the principal's problem. We relax the common prior assumption and assume that the distribution of the agents' types belongs to a commonly known ambiguity set. Assuming that the principal wants to maximize her worst-case expected payoff, we characterize optimal and Pareto robustly optimal mechanisms for three classes of ambiguity sets: (i) support-only ambiguity sets containing all distributions supported on a rectangle, (ii) Markov ambiguity sets containing all distributions supported on a rectangle whose mean values fall within another (smaller)

rectangle, and (iii) Markov ambiguity sets with independent types containing all distributions in Markov ambiguity sets under which agents' types are independent. Our results show that even in the presence of ambiguity, one can find a simple Pareto robustly optimal favored-agent mechanism. As opposed to the results with common prior assumption, the favored agent and the threshold of our Pareto robustly optimal mechanisms are independent of the verification costs.

Chapter 3 is a result of collaboration with my co-advisor, Dr. Nuh Aygün Dalkıran. We consider a freelancer who has limited working hours and can serve multiple parties by allocating his time accordingly. He has convex cost so that working an extra minute gets more costly as the allocated time for work gets high. On the other side of the market, several buyers are interested in the freelancer's services but have limited information regarding the service quality. Furthermore, no buyer can limit the freelancer regarding the contracts made with the other buyers. We characterize the equilibrium trades for this problem by assuming that there are at least two buyers, and the freelancer has private information regarding the quality of his service that can be either low or high. The buyers share a common prior regarding the quality of the service provided by the freelancer. The buyers have linear preferences for quality and compete through offering contracts that specify a quantity and a transfer. The freelancer observes the offers and chooses the contracts that maximize his payoff. The preferences of each type of the freelancer are quasilinear. In this context, we provide two distinct sufficiency conditions for the existence of a pure-strategy equilibrium. These conditions are also necessary for their respective equilibria. (i) At the no-trade point, the high-type freelancer is not willing to trade any amount of his time in exchange for a price equal to the average quality of the service. (ii) At any feasible trade point, the high-type freelancer is willing to trade any amount of his time in exchange for a price equal to the average quality of the service. If (i) holds, then the high-type freelancer does not trade in equilibrium, while the aggregate trade of the low-type depends on his preferences. In such equilibria, the buyers make zero profit from each traded contract. If (ii) holds, both types trade at the capacity, and there is cross-subsidization in equilibrium. In all of these equilibria, aggregate equilibrium trades are unique.

Chapter 4 is a result of collaborations with Dr. Rakesh Vohra and my advisor Dr. Mustafa Çelebi Pinar. We argue the usefulness of linear programming arguments in the derivation of prophet inequalities. Consider a gambler who is sequentially presented with the contents of many boxes, each containing a nonnegative reward. Each time a box is opened, the gambler must either claim the last unboxed reward or discard it and move on. The gambler wants to maximize his expected gains, whereas a prophet who can foresee all reward realizations can choose the highest reward in any game instance, hence gaining the expected value of the maximum reward. When the rewards are independently distributed, the optimal strategy of the gambler yields at least half of the optimal strategy of the prophet, which is called the prophet inequality. We contribute to the literature by re-establishing three different prophet inequality results by finding a feasible strategy for the gambler, which aligns with the prophet's optimal strategy in expectation to some desired level. (i) When the player is restricted to a polymatroid with nonnegative coefficients in its unique Minkowski sum of simplices, we prove the $\frac{1}{2}$ -prophet inequality. (ii) We show the $\frac{1}{n}$ -prophet inequality when there are n many boxes, the reward distributions can be dependent, and the gambler is restricted to an arbitrary polyhedron. (iii) When the gambler is subject to K many constraints, we obtain the $\frac{1}{K+1}$ -prophet inequality.

1.1.1 Statement of Originality

I certify that this dissertation is the result of my own work, where some parts are the results of collaborations with my advisor Dr. Mustafa Çelebi Pinar, my co-advisor Dr. Nuh Aygün Dalkıran, and my co-authors Dr. Çağıl Koçyiğit, Dr. Daniel Kuhn and Dr. Rakesh Vohra. No other person's work has been used without due acknowledgement.

Chapter 2

Distributionally Robust Optimal Allocation with Costly Verification

2.1 Introduction

Consider a principal ('she') who allocates a good to one of several agents without using monetary transfers. Each agent ('he') derives strictly positive utility from owning the good and has a private type, which reflects the value he creates for the principal if receiving the good. The principal is unaware of the agents' types but can verify any of them at a cost. Any verification will perfectly reveal the corresponding agent's type to the principal. The good is allocated based on the agents' self-declared types as well as the results of any verification performed. The principal aims to design an allocation mechanism that maximizes her payoff, *i.e.*, the value of allocation minus any costs of verification.

This generic mechanism design problem arises in many different contexts. For example, the rector of a university may have funding for a new faculty position and needs to

allocate it to one of the school's departments, the ministry of health may need to decide in which town to open up a new hospital, a venture capitalist may need to select a start-up business that should receive seed funding, the procurement manager of a manufacturing company may need to choose one of several suppliers, or a consulting company may need to identify a team that leads a new project. In all of these examples, the principal wishes to put the good into use where it best contributes to her organization or the society as a whole. Each agent desires the good and is likely to be well-informed about the value he will generate for the principal if he receives the good. In addition, monetary transfers may be inappropriate in all of the described situations, but the principal can collect information through costly investigation or audit.

Mechanism design problems of the above type are usually referred to as 'allocation with costly verification.' Ben-Porath et al. [3] describe the first formal model for their analysis and introduce the class of favored-agent mechanisms, which are attractive because of their simplicity and interpretability. As in most of the literature on mechanism design, [3] model the agents' types as independent random variables governed by a commonly known probability distribution, which allows them to prove that any mechanism that maximizes the principal's expected payoff is a randomization over favored-agent mechanisms. Any favored-agent mechanism is characterized by a favored agent and a threshold value, and it assigns the good to the favored agent without verification whenever the reported types of all other agents—adjusted for the costs of verification—fall below the given threshold. Otherwise, it allocates the good to any agent for which the reported type minus the cost of verification is maximal and verifies his reported type. The choice of the favored agent is predicated on the principal's prior beliefs about the agents' types. The favored agent receives the good without verification if no other agent claims to have a high enough type. Otherwise, the principal verifies the highest reported (adjusted) type and allocates the good to the respective agent. This mechanism is incentive compatible, that is, no agent has an incentive to misreport his true type; see Section 2.2 for more details.

The vast majority of the literature on allocation with costly verification (see, *e.g.*, [4, 5])

sustains the modeling assumptions of Ben-Porath et al. [3], thus assuming that the agents' types are independent random variables and that their distribution is common knowledge. In reality, however, it is often difficult to justify the precise knowledge of such a distribution. This prompts us to study allocation problems with costly verification under the more realistic assumption that the principal has only partial information about the distribution of the agents' types. Specifically, we assume that the distribution of the agents' types is unknown but belongs to a commonly known ambiguity set (*i.e.*, a family of multiple—perhaps infinitely many—distributions). In addition, we assume that the principal is ambiguity averse in the sense that she wishes to maximize her worst-case expected payoff in view of all distributions in the ambiguity set. Under these assumptions, the mechanism design problem at hand can be cast as a zero-sum game between the principal, who chooses a mechanism to allocate the good, and some fictitious adversary, who chooses the distribution of the agents' types from the ambiguity set in order to inflict maximum damage to the principal. Using techniques from distributionally robust optimization (see, *e.g.*, [6, 7]), we characterize optimal and Pareto robustly optimal mechanisms for three classes of ambiguity sets: (i) support-only ambiguity sets containing all distributions supported on a rectangle, (ii) Markov ambiguity sets containing all distributions supported on a rectangle whose mean values fall within another (smaller) rectangle, and (iii) Markov ambiguity sets with independent types containing all distributions in Markov ambiguity sets under which agents' types are independent. We emphasize that both support-only as well as Markov ambiguity sets contain distributions under which the agents' types are mutually dependent. Pareto robust optimality is an important solution concept in robust optimization (see [8]). In the distributionally robust context considered here, a mechanism is called Pareto robustly optimal if there is no other mechanism that generates a non-inferior expected payoff under every distribution in the ambiguity set and a strictly higher expected payoff under at least one distribution in the ambiguity set. Every Pareto robustly optimal solution is also robustly optimal, but the converse is not true. Mechanisms that fail to be Pareto robustly optimal would not be used by any rational agent.

For these three ambiguity sets, the contributions of this chapter can be summarized as

follows.

- (i) For support-only ambiguity sets, we first show that not every robustly optimal mechanism represents a randomization over favored-agent mechanisms. This result is unexpected in view of the classical theory on stochastic mechanism design, as Ben-Porath et al. [3] show that any optimal mechanism is a randomization over favored-agent mechanisms when the distribution of the agents' types is precisely known. We then construct an explicit favored-agent mechanism that is not only robustly optimal but also Pareto robustly optimal. This mechanism selects the favored agent from among those whose types have the highest possible lower bound, and it sets the threshold to this lower bound.
- (ii) For Markov ambiguity sets, we also construct an explicit favored-agent mechanism that is both robustly optimal as well as Pareto robustly optimal. This mechanism selects the favored agent from among those whose *expected* types have the highest possible lower bound, and it sets the threshold to the highest possible *actual* (not *expected*) type of the favored agent.
- (iii) For Markov ambiguity sets with independent types, we first prove that the principal's worst-case payoff coincides with the one under Markov ambiguity sets. Thus, the principal cannot increase her worst-case payoff with the additional knowledge of independence. We again construct a favored-agent mechanism that is optimal and Pareto robustly optimal. The favored agent of this mechanism chosen among those whose *expected* types have the highest possible lower bound (as in (ii)), and its threshold is set to this lower bound (unlike (ii)).

Our results show that favored-agent mechanisms continue to play an important role in allocation with costly verification even if the unrealistic assumption of a commonly known type distribution is abandoned. In addition, they suggest that robust optimality alone may not be a sufficiently distinctive criterion to single out practically useful mechanisms under distributional ambiguity. However, our results also show that among possibly infinitely

many robustly optimal mechanisms one can always find a simple and interpretable Pareto robustly optimal favored-agent mechanism. Unlike Ben-Porath et al. [3], the favored agent as well as the threshold of our Pareto robustly optimal mechanisms are *independent* of the verification costs.

Literature review. The first treatise of allocation with costly verification is due to Townsend [9], who studies a principal-agent model with monetary transfers involving a single agent. Ben-Porath et al. [3] extend this model to multiple agents but rule out the possibility of monetary transfers. Their seminal work has inspired considerable follow-up research in economics. For example, Mylovanov & Zapechelnyuk [5] study a variant of the problem where verification is costless but the principal can impose only limited penalties and only partially recover the good when agents misreport their types. Li [4] accounts both for costly verification and for limited penalties, thereby unifying the models in [3] and [5]. Chua et al. [10] further extend the model in [3] to multiple homogeneous goods, assuming that each agent can receive at most one good. Bayrak et al. [11] spearhead the study of allocation with costly verification under distributional ambiguity. However, for reasons of computational tractability, they focus on ambiguity sets that contain only two discrete distributions. In this chapter, we investigate ambiguity sets that contain infinitely many (not necessarily discrete) type distributions characterized by support and moment constraints, and we derive robustly as well as Pareto robustly optimal mechanisms in closed form.

This chapter also contributes to the growing literature on (distributionally) robust mechanism design. Note that any mechanism design problem is inherently affected by uncertainty due to the private information held by different agents. The vast majority of the extant mechanism design literature models uncertainty through random variables that are governed by a commonly known probability distribution. The robust mechanism design literature, on the other hand, explicitly accounts for (non-stochastic) distributional uncertainty and seeks mechanisms that maximize the worst-case payoff, minimize the worst-case regret or minimize the worst-case cost in view of all distributions consistent with the

information available. Robust mechanism design problems have recently emerged in different contexts such as pricing (see, *e.g.*, [12, 13, 14, 15, 16, 17]), auction design (see, *e.g.*, [18, 19, 20, 21]) or contracting (see, *e.g.*, [22]). The literature in this area is too vast to discuss all contributions in detail. To our best knowledge, however, this chapter is the first one that derives closed-form optimal mechanisms for the allocation problem with costly verification under distributional ambiguity.

The remainder of this chapter is structured as follows. Section 2.2 introduces our model and establishes several preliminary results. Sections 2.3, 2.4, and 2.5 solve the proposed mechanism design problem for support-only, Markov and Markov with independence ambiguity sets, respectively.

Notation. For any $\mathbf{t} \in \mathbb{R}^I$, we denote by t_i the i^{th} component and by \mathbf{t}_{-i} the subvector of \mathbf{t} without t_i . The indicator function of a logical expression E is defined as $\mathbb{1}_E = 1$ if E is true and as $\mathbb{1}_E = 0$ otherwise. For any Borel sets $\mathcal{S} \subseteq \mathbb{R}^n$ and $\mathcal{D} \subseteq \mathbb{R}^m$, we use $\mathcal{P}_0(\mathcal{S})$ and $\mathcal{L}(\mathcal{S}, \mathcal{D})$ to denote the family of all probability distributions on \mathcal{S} and the set of all bounded Borel-measurable functions from \mathcal{S} to \mathcal{D} , respectively. Random variables are designated by symbols with tildes (*e.g.*, $\tilde{\mathbf{t}}$), and their realizations are denoted by the same symbols without tildes (*e.g.*, \mathbf{t}).

2.2 Problem Statement and Preliminaries

A principal wishes to allocate a single good to one of $I \geq 2$ agents. Each agent $i \in \mathcal{I} = \{1, 2, \dots, I\}$ derives a strictly positive deterministic benefit from receiving the good and also uses it to generate a value $t_i \in \mathcal{T}_i = [\underline{t}_i, \bar{t}_i]$ for the principal, where $0 \leq \underline{t}_i < \bar{t}_i < \infty$. We henceforth refer to t_i as agent i 's type, and we assume that t_i is privately known to agent i but unknown to the principal and the other agents. Thus, the principal perceives the vector $\tilde{\mathbf{t}} = (\tilde{t}_1, \tilde{t}_2, \dots, \tilde{t}_I)$ of all agents' types as a random vector governed by some

probability distribution \mathbb{P}_0 on the type space $\mathcal{T} = \prod_{i \in \mathcal{I}} \mathcal{T}_i$. However, the principal can inspect agent i 's type at cost $c_i > 0$, and the inspection perfectly reveals t_i . In contrast to much of the existing literature on mechanism design, we assume here that neither the principal nor the agents know \mathbb{P}_0 . Instead, they are only aware that \mathbb{P}_0 belongs to some commonly known ambiguity set $\mathcal{P} \subseteq \mathcal{P}_0(\mathcal{T})$. On this basis, the principal aims to design a mechanism for allocating the good. A mechanism is an extensive-form game between the principal and the agents, where the principal commits in advance to her strategy (for a formal definition of extensive form games, see, *e.g.*, [23]). Such a mechanism may contain multiple stages of cheap talk statements by the agents, while the principal's actions include the decisions on whether to inspect certain agents and how to allocate the good. Monetary transfers are not allowed, *i.e.*, the agents and the principal cannot exchange money at any time.

Given any mechanism represented as an extensive form game, we denote by \mathcal{H}_i the family of all information sets of agent i and by $\mathcal{A}(h_i)$ the actions available to agent i at the nodes in information set $h_i \in \mathcal{H}_i$. All agents select their actions strategically in view of their individual preferences and the available information. In particular, agent i 's actions depend on his type t_i . Thus, we model any (mixed) strategy of agent i as a function $s_i \in \mathcal{L}(\mathcal{T}_i, \prod_{h_i \in \mathcal{H}_i} \mathcal{P}_0(\mathcal{A}(h_i)))$ that maps each of his possible types to a complete contingency plan $a_i \in \prod_{h_i \in \mathcal{H}_i} \mathcal{P}_0(\mathcal{A}(h_i))$, which represents a probability distribution over the actions available to agent i for all information sets $h_i \in \mathcal{H}_i$. In the following we denote by $\text{prob}_i(a_i; \mathbf{t}, \mathbf{a}_{-i})$ the probability that agent $i \in \mathcal{I}$ receives the good under the principal's mechanism if the agents have types \mathbf{t} and play the contingency plans $\mathbf{a} = (a_1, a_2, \dots, a_I)$. We also restrict attention to mechanisms that admit an ex-post Nash equilibrium.

Definition 1 (Ex-Post Nash Equilibrium). *An I -tuple $\mathbf{s} = (s_1, s_2, \dots, s_I)$ of mixed strategies $s_i \in \mathcal{L}(\mathcal{T}_i, \prod_{h_i \in \mathcal{H}_i} \mathcal{P}_0(\mathcal{A}(h_i)))$, $i \in \mathcal{I}$, is called an ex-post Nash equilibrium if*

$$\begin{aligned} \text{prob}_i(s_i(t_i); \mathbf{t}, \mathbf{s}_{-i}(\mathbf{t}_{-i})) &\geq \text{prob}_i(a_i; \mathbf{t}, \mathbf{s}_{-i}(\mathbf{t}_{-i})) \\ &\forall i \in \mathcal{I}, \forall \mathbf{t} \in \mathcal{T}, \forall a_i \in \prod_{h_i \in \mathcal{H}_i} \mathcal{P}_0(\mathcal{A}(h_i)). \end{aligned}$$

Recall that all agents assign a strictly positive deterministic value to the good, and therefore the expected utility of agent i conditional on $\tilde{\mathbf{t}} = \mathbf{t}$ is proportional to $\text{prob}_i(a_i; \mathbf{t}, \mathbf{a}_{-i})$. Under an ex-post Nash equilibrium, each agent i maximizes this probability simultaneously for all type scenarios $\mathbf{t} \in \mathcal{T}$. Hence, it is clear that insisting on the existence of an ex-post Nash equilibrium restricts the family of mechanisms to be considered. Note that Ben-Porath et al. [3] study the larger class of mechanisms that admit a Bayesian Nash equilibrium. However, these mechanisms generically depend on the type distribution \mathbb{P}_0 and can therefore not be implemented by a principle who lacks knowledge of \mathbb{P}_0 . It is therefore natural to restrict attention to mechanisms that admit ex-post Nash equilibria, which remain well-defined in the face of distributional ambiguity. We further assume from now on that the principal is ambiguity averse in the sense that she wishes to maximize her worst-case expected payoff in view of all distributions in the ambiguity set \mathcal{P} .

The class of all mechanisms that admit an ex-post Nash equilibrium is vast. An important subclass is the family of all truthful direct mechanisms. A direct mechanism (\mathbf{p}, \mathbf{q}) consists of two I -tuples $\mathbf{p} = (p_1, p_2, \dots, p_I)$ and $\mathbf{q} = (q_1, q_2, \dots, q_I)$ of allocation functions $p_i, q_i \in \mathcal{L}(\mathcal{T}, [0, 1])$, $i \in \mathcal{I}$. Any direct mechanism (\mathbf{p}, \mathbf{q}) is implemented as follows. First, the principal announces \mathbf{p} and \mathbf{q} , and then she collects a bid $t'_i \in \mathcal{T}_i$ from each agent $i \in \mathcal{I}$. Next, the principal implements randomized allocation and inspection decisions. Specifically, $p_i(\mathbf{t}')$ represents the total probability that agent i receives the good, while $q_i(\mathbf{t}')$ represents the probability that agent i receives the good *and* is inspected. If the inspection reveals that agent i has misreported his type, the principal penalizes the agent by repossessing the good. Any direct mechanism (\mathbf{p}, \mathbf{q}) must satisfy the feasibility conditions

$$q_i(\mathbf{t}') \leq p_i(\mathbf{t}') \quad \forall i \in \mathcal{I} \quad \text{and} \quad \sum_{i \in \mathcal{I}} p_i(\mathbf{t}') \leq 1 \quad \forall \mathbf{t}' \in \mathcal{T}. \quad (\text{FC})$$

The first inequality in (FC) holds because only agents who receive the good may undergo an inspection. The second inequality in (FC) ensures that the principal allocates the good at most once.

A direct mechanism (\mathbf{p}, \mathbf{q}) is called truthful if it is optimal for each agent i to report his true type $t'_i = t_i$. Thus, (\mathbf{p}, \mathbf{q}) is truthful if and only if it satisfies the incentive compatibility constraints

$$p_i(\mathbf{t}) \geq p_i(t'_i, \mathbf{t}_{-i}) - q_i(t'_i, \mathbf{t}_{-i}) \quad \forall i \in \mathcal{I}, \forall t'_i \in \mathcal{T}_i, \forall \mathbf{t} \in \mathcal{T}, \quad (\text{IC})$$

which ensure that if all other agents report their true types \mathbf{t}_{-i} , then the probability $p_i(\mathbf{t})$ of agent i receiving the good if he reports his true type t_i exceeds the probability $p_i(t'_i, \mathbf{t}_{-i}) - q_i(t'_i, \mathbf{t}_{-i})$ of agent i receiving the good if he misreports his type as $t'_i \neq t_i$. By leveraging a variant of the Revelation Principle detailed in Ben-Porath et al. [3], one can show that for any mechanism that admits an ex-post Nash equilibrium there exists an equivalent truthful direct mechanism that duplicates or improves the principal's worst-case expected payoff; see the online appendix of [3] for details. Without loss of generality, the principal may thus focus on truthful direct mechanisms, which greatly simplifies the problem of finding an optimal mechanism. Consequently, the principal's mechanism design problem can be formalized as the following distributionally robust optimization problem.

$$\begin{aligned} z^* &= \sup_{\mathbf{p}, \mathbf{q}} \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} (p_i(\tilde{\mathbf{t}}) \tilde{t}_i - q_i(\tilde{\mathbf{t}}) c_i) \right] \\ \text{s.t.} \quad & p_i, q_i \in \mathcal{L}(\mathcal{T}, [0, 1]) \quad \forall i \in \mathcal{I} \\ & (\text{IC}), (\text{FC}) \end{aligned} \quad (\text{MDP})$$

From now on, we will use the shorthand \mathcal{X} to denote the set of all (\mathbf{p}, \mathbf{q}) feasible in (MDP). Note that the feasible set \mathcal{X} does not rely on the ambiguity set and remains the same for all $\mathcal{P} \subseteq \mathcal{P}_0(\mathcal{T})$.

In the remainder we will demonstrate that (MDP) often admits multiple optimal solutions. While different optimal mechanisms generate the same expected profit in the worst case, they may offer dramatically different expected profits under generic non-worst-case

distributions. This observation prompts us to seek mechanisms that are not only worst-case optimal but perform also well under *all* type distributions in the ambiguity set \mathcal{P} . More precisely, we hope to identify a worst-case optimal mechanism for which there exists no other feasible mechanism that generates at least the same expected payoff under *every* distribution in \mathcal{P} and a higher expected payoff under *at least one* distribution in \mathcal{P} . A mechanism with this property is called *Pareto robustly optimal*. This terminology is borrowed from the theory of Pareto efficiency in classical robust optimization (see [8] for details).

Definition 2 (Pareto Robust Optimality). *We say that a mechanism $(\mathbf{p}', \mathbf{q}')$ that is feasible in (MDP) weakly Pareto robustly dominates another feasible mechanism (\mathbf{p}, \mathbf{q}) if*

$$\mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} (p'_i(\tilde{\mathbf{t}}) \tilde{t}_i - q'_i(\tilde{\mathbf{t}}) c_i) \right] \geq \mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} (p_i(\tilde{\mathbf{t}}) \tilde{t}_i - q_i(\tilde{\mathbf{t}}) c_i) \right] \quad \forall \mathbb{P} \in \mathcal{P}. \quad (2.1)$$

If (2.1) holds and the inequality is strict for at least one $\mathbb{P} \in \mathcal{P}$, then we say that $(\mathbf{p}', \mathbf{q}')$ Pareto robustly dominates (\mathbf{p}, \mathbf{q}) . A mechanism (\mathbf{p}, \mathbf{q}) that is optimal in (MDP) is called Pareto robustly optimal if there exists no other feasible mechanism $(\mathbf{p}', \mathbf{q}')$ that Pareto robustly dominates (\mathbf{p}, \mathbf{q}) .

Note that any mechanism that weakly Pareto robustly dominates an optimal mechanism is also optimal in (MDP). Moreover, a Pareto robustly optimal mechanism always exists. However, there may not exist any mechanism that Pareto robustly dominates all other feasible mechanisms.

We now define the notion of a favored-agent mechanism, which was first introduced in Ben-Porath et al. [3].

Definition 3 (Favored-Agent Mechanism). *A mechanism (\mathbf{p}, \mathbf{q}) is a favored-agent mechanism if there exists a favored agent $i^* \in \mathcal{I}$ and a threshold value $\nu^* \in \mathbb{R}$ such that the following hold.*

- (i) If $\max_{i \neq i^*} t_i - c_i < \nu^*$, then $p_{i^*}(\mathbf{t}) = 1$, $q_{i^*}(\mathbf{t}) = 0$ and $p_i(\mathbf{t}) = q_i(\mathbf{t}) = 0$ for all $i \neq i^*$.
- (ii) If $\max_{i \neq i^*} t_i - c_i > \nu^*$, then $p_{i'}(\mathbf{t}) = q_{i'}(\mathbf{t}) = 1$ for some $i' \in \arg \max_{i \in \mathcal{I}} (t_i - c_i)$ and $p_i(\mathbf{t}) = q_i(\mathbf{t}) = 0$ for all $i \neq i'$.

If $\max_{i \neq i^*} t_i - c_i = \nu^*$, then we are free to define $(\mathbf{p}(\mathbf{t}), \mathbf{q}(\mathbf{t}))$ either as in (i) or as in (ii).

Intuitively, if t_i is smaller than the adjusted cost of inspection $c_i + \nu^*$ for every agent $i \neq i^*$, then any favored-agent mechanism allocates the good to the favored agent i^* without inspection. If there exists an agent $i \neq i^*$ whose type t_i exceeds the adjusted cost of inspection $c_i + \nu^*$, then the favored-agent mechanism allocates the good to an agent i' with highest net payoff $t_{i'} - c_{i'}$, and this agent is inspected. Note that in case (ii) the good can also be allocated to the favored agent.

A favored-agent mechanism is uniquely determined by a favored agent i^* , a threshold value ν^* and two tie-breaking rules. The first tie-breaking rule determines the winning agent in case (ii) when $\arg \max_{i \in \mathcal{I}} (t_i - c_i)$ is not a singleton. From now on we will always use the lexicographic tie-breaking rule in this case, which sets $i' = \min \arg \max_{i \in \mathcal{I}} (t_i - c_i)$. The second tie-breaking rule determines whether $(\mathbf{p}(\mathbf{t}), \mathbf{q}(\mathbf{t}))$ should be constructed as in case (i) or as in case (ii) when $\max_{i \neq i^*} t_i - c_i = \nu^*$. From now on we say that a favored-agent mechanism is of type (i) if $(\mathbf{p}(\mathbf{t}), \mathbf{q}(\mathbf{t}))$ is always defined as in (i) and that it is of type (ii) if $(\mathbf{p}(\mathbf{t}), \mathbf{q}(\mathbf{t}))$ is always defined as in (ii) in case of a tie. Note that both tie-breaking rules are irrelevant in the Bayesian setting considered in Ben-Porath et al. [3], but they are relevant for us because the ambiguity sets \mathcal{P} to be studied below contain discrete distributions, under which ties have a strictly positive probability.

All favored-agent mechanisms are feasible in (MDP), see Remark 1 in [3]. In particular, they are incentive compatible, that is, the agents have no incentive to misreport their types. To see this, recall that under a favored-agent mechanism the winning agent receives the good with probability one, and the losing agents receive the good with probability zero.

Thus, if an agent wins by truthful bidding, he cannot increase his chances of receiving the good by lying about his type. If an agent loses by truthful bidding, on the other hand, he has certainly no incentive to lower his bid t_i because the chances of receiving the good are non-decreasing in t_i . Increasing his bid t_i may earn him the good provided that $t_i - c_i$ attains the maximum of $t_{i'} - c_{i'}$ over $i' \in \mathcal{I}$. However, in this case the agent's type is inspected with probability one. Hence, the lie will be detected and the good will be repossessed. This shows that no agent benefits from lying under a favored-agent mechanism.

If $\mathcal{P} = \{\mathbb{P}_0\}$ is a singleton, the agents' types are independent under \mathbb{P}_0 , and \mathbb{P}_0 has an everywhere positive density on \mathcal{T} , then problem (MDP) is solved by a favored-agent mechanism, see Ben-Porath et al. [3, Theorem 1]. The favored-agent mechanism with favored agent i and threshold ν_i yields an expected payoff of

$$\begin{aligned} & \mathbb{E}_{\mathbb{P}_0} [\tilde{t}_i \mathbb{1}_{\tilde{y}_i \leq \nu_i} + \max \{ \tilde{t}_i - c_i, \tilde{y}_i \} \mathbb{1}_{\tilde{y}_i \geq \nu_i}] \\ &= \int_{-\infty}^{\nu_i} \mathbb{E}_{\mathbb{P}_0} [\tilde{t}_i] \rho_i(y_i) dy_i + \int_{\nu_i}^{\infty} \mathbb{E}_{\mathbb{P}_0} [\max \{ \tilde{t}_i - c_i, y_i \}] \rho_i(y_i) dy_i, \end{aligned}$$

where $\rho_i(y_i)$ denotes the probability density function of the random variable $\tilde{y}_i = \max_{j \neq i} \tilde{t}_j - c_j$ with respect to \mathbb{P}_0 , which is independent of \tilde{t}_i under \mathbb{P}_0 . The threshold value ν_i^* that maximizes the above expression for a fixed favored agent i thus solves the first-order optimality condition

$$\mathbb{E}_{\mathbb{P}_0} [\tilde{t}_i] = \mathbb{E}_{\mathbb{P}_0} [\max \{ \tilde{t}_i - c_i, \nu_i \}]. \quad (2.2)$$

Note that ν_i^* is unique because the right hand side of (2.2) strictly increases in ν_i on the domain of interest, see [3, Theorem 2]. One can further prove that within the finite class of favored-agent mechanisms with optimal thresholds, the ones with the highest threshold are optimal. More specifically, any favored-agent mechanism with favored agent $i^* \in \arg \max_{i \in \mathcal{I}} \nu_i^*$ and threshold $\nu^* = \max_{i \in \mathcal{I}} \nu_i^*$ is optimal within the class of favored-agent mechanisms, see [3, Theorem 3]. Hence, any such mechanism must be optimal in (MDP). Finally, one can also show that for mutually distinct cost coefficients $c_i, i \in \mathcal{I}$, the optimal

favored-agent mechanism is unique.

In the remainder of the chapter, we will address instances of the mechanism design problem (MDP) where \mathcal{P} is *not* a singleton, and we will prove that favored-agent mechanisms remain optimal. Under distributional ambiguity, however, the construction of i^* and ν^* described above is no longer well-defined because it depends on a particular choice of the probability distribution of $\tilde{\mathbf{t}}$. We will show that if \mathcal{P} is not a singleton, then there may be infinitely many optimal favored-agent mechanisms with different thresholds ν^* . In this situation, it is expedient to look for Pareto robustly optimal favored-agent mechanisms.

2.3 Support-Only Ambiguity Sets

We now investigate the mechanism design problem (MDP) under the assumption that $\mathcal{P} = \mathcal{P}_0(\mathcal{T})$ is the support-only ambiguity set that contains all distributions supported on the type space \mathcal{T} . As \mathcal{P} contains all Dirac point distributions concentrating unit mass at any $\mathbf{t} \in \mathcal{T}$, the worst-case expected payoff over all distributions $\mathbb{P} \in \mathcal{P}$ simplifies to the worst-case payoff over all type profiles $\mathbf{t} \in \mathcal{T}$, and thus it is easy to verify that problem (MDP) simplifies to

$$\begin{aligned}
 z^* &= \sup_{\mathbf{p}, \mathbf{q}} \inf_{\mathbf{t} \in \mathcal{T}} \sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) \\
 \text{s.t. } & p_i, q_i \in \mathcal{L}(\mathcal{T}, [0, 1]) \quad \forall i \in \mathcal{I} \\
 & \text{(IC), (FC)}.
 \end{aligned} \tag{2.3}$$

Similarly, it is easy to verify that an optimal mechanism $(\mathbf{p}^*, \mathbf{q}^*)$ for problem (2.3) is Pareto robustly optimal if there exists no other feasible mechanism (\mathbf{p}, \mathbf{q}) with

$$\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) \geq \sum_{i \in \mathcal{I}} (p_i^*(\mathbf{t})t_i - q_i^*(\mathbf{t})c_i) \quad \forall \mathbf{t} \in \mathcal{T},$$

where the inequality is strict for at least one type profile $\mathbf{t} \in \mathcal{T}$. If the principal knew the agents' types ex ante, she could simply allocate the good to the agent with the highest type and would not have to spend money on inspecting anyone. One can therefore show that the optimal value z^* of problem (2.3) is upper bounded by $\inf_{\mathbf{t} \in \mathcal{T}} \max_{i \in \mathcal{I}} t_i = \max_{i \in \mathcal{I}} \underline{t}_i$. The following proposition formally establishes this upper bound and shows that it is indeed attained by an admissible mechanism.

Proposition 1. *The optimal value of (2.3) is given by $z^* = \max_{i \in \mathcal{I}} \underline{t}_i$.*

Proof. Relaxing the incentive compatibility constraints and the first inequality in (FC) yields

$$\begin{aligned}
z^* &\leq \sup_{\mathbf{p}, \mathbf{q}} \inf_{\mathbf{t} \in \mathcal{T}} \sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) \\
&\text{s.t. } p_i, q_i \in \mathcal{L}(\mathcal{T}, [0, 1]) \quad \forall i \in \mathcal{I} \\
&\quad \sum_{i \in \mathcal{I}} p_i(\mathbf{t}) \leq 1 \quad \forall \mathbf{t} \in \mathcal{T} \\
&= \sup_{\mathbf{p}} \inf_{\mathbf{t} \in \mathcal{T}} \sum_{i \in \mathcal{I}} p_i(\mathbf{t})t_i \\
&\text{s.t. } p_i \in \mathcal{L}(\mathcal{T}, [0, 1]) \quad \forall i \in \mathcal{I}, \quad \sum_{i \in \mathcal{I}} p_i(\mathbf{t}) \leq 1 \quad \forall \mathbf{t} \in \mathcal{T},
\end{aligned}$$

where the equality holds because in the relaxed problem it is optimal to set $q_i(\mathbf{t}) = 0$ for all $i \in \mathcal{I}$ and $\mathbf{t} \in \mathcal{T}$. As the resulting maximization problem over \mathbf{p} is separable with respect to $\mathbf{t} \in \mathcal{T}$, it is optimal to allocate the good in each scenario $\mathbf{t} \in \mathcal{T}$ —with probability one—to an agent with maximal type. Therefore, z^* is bounded above by $\inf_{\mathbf{t} \in \mathcal{T}} \max_{i \in \mathcal{I}} t_i = \max_{i \in \mathcal{I}} \underline{t}_i$. However, this bound is attained by a mechanism that allocates the good to an agent $i' \in \arg \max_{i \in \mathcal{I}} \underline{t}_i$ irrespective of $\mathbf{t} \in \mathcal{T}$ and never inspects anyone's type. Since this mechanism is feasible, the claim follows. \square

The next theorem shows that there are infinitely many optimal favored-agent mechanisms that attain the optimal value $z^* = \max_{i \in \mathcal{I}} \underline{t}_i$ of problem (2.3).

Theorem 1. *Any favored-agent mechanism with favored agent $i^* \in \arg \max_{i \in \mathcal{I}} t_i$ and threshold value $\nu^* \geq \max_{i \in \mathcal{I}} \underline{t}_i$ is optimal in problem (2.3).*

Proof. Select an arbitrary favored-agent mechanism with $i^* \in \arg \max_{i \in \mathcal{I}} t_i$ and $\nu^* \geq \max_{i \in \mathcal{I}} \underline{t}_i$. Recall first that this mechanism is feasible in (2.3). Next, we will show that this mechanism attains a worst-case payoff that is at least as large as $\max_{i \in \mathcal{I}} \underline{t}_i$, which implies via Proposition 1 that it is in fact optimal in (2.3). To this end, fix an arbitrary type profile $\mathbf{t} \in \mathcal{T}$. If $\max_{i \neq i^*} t_i - c_i < \nu^*$, then condition (i) in Definition 3 implies that the principal's payoff amounts to $t_{i^*} \geq \max_{i \in \mathcal{I}} \underline{t}_i$, where the inequality follows from the selection of i^* . If $\max_{i \neq i^*} t_i - c_i > \nu^*$, then condition (ii) in Definition 3 implies that the principal's payoff amounts to $\max_{i \in \mathcal{I}} t_i - c_i > \nu^* \geq \max_{i \in \mathcal{I}} \underline{t}_i$, where the second inequality follows from the selection of ν^* . If $\max_{i \neq i^*} t_i - c_i = \nu^*$, then the allocation functions are defined either as in condition (i) or as in condition (ii) of Definition 3. Thus, the principal's payoff amounts either to t_{i^*} or to $\max_{i \in \mathcal{I}} t_i - c_i \geq \nu^*$, respectively, and is therefore again non-inferior to $\max_{i \in \mathcal{I}} \underline{t}_i$. In summary, we have shown that the principal's payoff is non-inferior to $z^* = \max_{i \in \mathcal{I}} \underline{t}_i$ in all three cases. As scenario $\mathbf{t} \in \mathcal{T}$ was chosen arbitrarily, this reasoning implies that the principal's worst-case payoff is also non-inferior to z^* . The favored-agent mechanism at hand is therefore optimal in (2.3) by virtue of Proposition 1. \square

As the mechanism design problem (2.3) constitutes a convex program, any convex combination of optimal favored-agent mechanisms gives rise to yet another optimal mechanism. However, problem (2.3) also admits optimal mechanisms that can neither be interpreted as favored-agent mechanisms nor as convex combinations of favored-agent mechanisms. To see this, consider any favored-agent mechanism (\mathbf{p}, \mathbf{q}) with favored agent $i^* \in \arg \max_{i \in \mathcal{I}} \underline{t}_i$ and threshold value $\nu^* \in \mathbb{R}$ satisfying $\nu^* \geq \max_{i \in \mathcal{I}} \underline{t}_i$ and $\nu^* > \max_{i \in \mathcal{I}} \bar{t}_i - c_i$. By Theorem 1, this mechanism is optimal in problem (2.3). The second condition on ν^* implies that this mechanism always allocates the good to the favored agent without inspection for every $\mathbf{t} \in \mathcal{T}$ (case (i) always prevails). Next, construct $\hat{\mathbf{t}} \in \mathcal{T}$

through $\hat{t}_i = \underline{t}_i$ for all $i \neq i^*$ and $\hat{t}_{i^*} = \bar{t}_{i^*}$, and note that $\hat{\mathbf{t}} \neq \underline{\mathbf{t}}$ because $\underline{t}_{i^*} < \bar{t}_{i^*}$. Finally, introduce another mechanism $(\mathbf{p}, \mathbf{q}')$, where \mathbf{q}' is defined through $q'_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $\mathbf{t} \in \mathcal{T}$ and $i \neq i^*$ and

$$q'_{i^*}(\mathbf{t}) = \begin{cases} \min\{1, (\bar{t}_{i^*} - \underline{t}_{i^*})/c_{i^*}\} & \text{if } \mathbf{t} = \hat{\mathbf{t}}, \\ q_{i^*}(\mathbf{t}) & \text{if } \mathbf{t} \in \mathcal{T} \setminus \{\hat{\mathbf{t}}\}. \end{cases}$$

One readily verifies that $(\mathbf{p}, \mathbf{q}')$ is feasible in (2.3). Indeed, as $(\mathbf{p}, \mathbf{q}')$ differs from (\mathbf{p}, \mathbf{q}) only in scenario $\hat{\mathbf{t}}$ and as (\mathbf{p}, \mathbf{q}) is feasible, it suffices to check the feasibility of $(\mathbf{p}, \mathbf{q}')$ in scenario $\hat{\mathbf{t}}$. For example, the modified allocation rule \mathbf{q}' is valued in $[0, 1]$ and $(\mathbf{p}, \mathbf{q}')$ satisfies (FC) because

$$0 \leq q'_{i^*}(\hat{\mathbf{t}}) \leq 1 = p_{i^*}(\hat{\mathbf{t}}),$$

where the equality holds because the favored-agent mechanism (\mathbf{p}, \mathbf{q}) allocates the good with certainty to agent i^* . Similarly, the modified mechanism $(\mathbf{p}, \mathbf{q}')$ satisfies (IC) because

$$p_{i^*}(t_{i^*}, \hat{\mathbf{t}}_{-i^*}) = 1 \geq p_{i^*}(\hat{\mathbf{t}}) - q'_{i^*}(\hat{\mathbf{t}}) \quad \forall t_{i^*} \in \mathcal{T}_{i^*}.$$

In summary, we have thus shown that the mechanism $(\mathbf{p}, \mathbf{q}')$ is feasible in (2.3). To show that it is also optimal, recall that (\mathbf{p}, \mathbf{q}) is optimal with worst-case payoff $\max_{i \in \mathcal{I}} \underline{t}_i$ and that $(\mathbf{p}, \mathbf{q}')$ differs from (\mathbf{p}, \mathbf{q}) only in scenario $\hat{\mathbf{t}}$. The principal's payoff in scenario $\hat{\mathbf{t}}$ amounts to

$$p_{i^*}(\hat{\mathbf{t}})\hat{t}_{i^*} - q'_{i^*}(\hat{\mathbf{t}})c_{i^*} = \hat{t}_{i^*} - q'_{i^*}(\hat{\mathbf{t}})c_{i^*} \geq \hat{t}_{i^*} - \frac{\hat{t}_{i^*} - \underline{t}_{i^*}}{c_{i^*}}c_{i^*} = \underline{t}_{i^*} = \max_{i \in \mathcal{I}} \underline{t}_i,$$

where the inequality follows from the definition of $q'_{i^*}(\hat{\mathbf{t}})$. Thus, the worst-case payoff of $(\mathbf{p}, \mathbf{q}')$ amounts to $\max_{i \in \mathcal{I}} \underline{t}_i$, and $(\mathbf{p}, \mathbf{q}')$ is indeed optimal in (2.3). However, $(\mathbf{p}, \mathbf{q}')$ is *not* a favored-agent mechanism for otherwise $q'_{i^*}(\hat{\mathbf{t}})$ would have to vanish; see Definition 3. In addition, note that $p_{i^*}(\hat{\mathbf{t}}) - q'_{i^*}(\hat{\mathbf{t}}) < 1$ whereas $p_{i^*}(t_{i^*}, \hat{\mathbf{t}}_{-i^*}) - q'_{i^*}(t_{i^*}, \hat{\mathbf{t}}_{-i^*}) = 1$ for all $t_{i^*} \neq \hat{t}_{i^*}$. This implies via Lemma 1 below that $(\mathbf{p}, \mathbf{q}')$ is also *not* a convex combination of favored-agent mechanisms.

Lemma 1. *If a mechanism (\mathbf{p}, \mathbf{q}) is a convex combination of favored-agent mechanisms, then the function $p_i(t_i, \mathbf{t}_{-i}) - q_i(t_i, \mathbf{t}_{-i})$ is constant in $t_i \in \mathcal{T}_i$ for any fixed $i \in \mathcal{I}$ and $\mathbf{t}_{-i} \in \mathcal{T}_{-i}$.*

Proof. Assume first that (\mathbf{p}, \mathbf{q}) is a favored-agent mechanism with favored agent $i^* \in \mathcal{I}$ and threshold value $\nu^* \in \mathbb{R}$. Next, fix any agent $i \in \mathcal{I}$ and any type profile $\mathbf{t}_{-i} \in \mathcal{T}_{-i}$. If $i \neq i^*$, then we have either $p_i(t_i, \mathbf{t}_{-i}) = q_i(t_i, \mathbf{t}_{-i}) = 1$ or $p_i(t_i, \mathbf{t}_{-i}) = q_i(t_i, \mathbf{t}_{-i}) = 0$ for all $t_i \in \mathcal{T}_i$. This implies that $p_i(t_i, \mathbf{t}_{-i}) - q_i(t_i, \mathbf{t}_{-i}) = 0$ is constant in $t_i \in \mathcal{T}_i$. If $i = i^*$, then the fixed type profile \mathbf{t}_{-i^*} uniquely determines whether the allocations are constructed as in case (i) or as in case (ii) of Definition 3. In case (i) we have $p_{i^*}(t_{i^*}, \mathbf{t}_{-i^*}) = 1$ and $q_{i^*}(t_{i^*}, \mathbf{t}_{-i^*}) = 0$ for all $t_{i^*} \in \mathcal{T}_{i^*}$, and thus $p_{i^*}(t_{i^*}, \mathbf{t}_{-i^*}) - q_{i^*}(t_{i^*}, \mathbf{t}_{-i^*}) = 1$ is constant in $t_{i^*} \in \mathcal{T}_{i^*}$. In case (ii) we have either $p_{i^*}(t_{i^*}, \mathbf{t}_{-i^*}) = 1$ and $q_{i^*}(t_{i^*}, \mathbf{t}_{-i^*}) = 1$ or $p_{i^*}(t_{i^*}, \mathbf{t}_{-i^*}) = 0$ and $q_{i^*}(t_{i^*}, \mathbf{t}_{-i^*}) = 0$, and thus $p_{i^*}(t_{i^*}, \mathbf{t}_{-i^*}) - q_{i^*}(t_{i^*}, \mathbf{t}_{-i^*}) = 0$ is again constant in $t_{i^*} \in \mathcal{T}_{i^*}$. This establishes the claim for any favored-agent mechanism (\mathbf{p}, \mathbf{q}) . Assume now that $(\mathbf{p}, \mathbf{q}) = \sum_{k \in \mathcal{K}} \pi_k(\mathbf{p}^k, \mathbf{q}^k)$ is a convex combination of favored-agent mechanisms $(\mathbf{p}^k, \mathbf{q}^k)$, $k \in \mathcal{K} = \{1, \dots, K\}$. Next, fix any $i \in \mathcal{I}$ and $\mathbf{t}_{-i} \in \mathcal{T}_{-i}$. From the first part of the proof we know that $p_i^k(t_i, \mathbf{t}_{-i}) - q_i^k(t_i, \mathbf{t}_{-i})$ is constant in $t_i \in \mathcal{T}_i$ for each $k \in \mathcal{K}$, and therefore $p_i(t_i, \mathbf{t}_{-i}) - q_i(t_i, \mathbf{t}_{-i})$ is also constant in $t_i \in \mathcal{T}_i$. Similar arguments apply when (\mathbf{p}, \mathbf{q}) represents a convex combination of infinitely many favored-agent mechanisms. \square

The above reasoning implies that the robust mechanism design problem (2.3) admits infinitely many optimal solutions. Some of these solutions represent favored-agent mechanisms while some represent convex combinations of favored-agent mechanisms, and yet some others are different types of mechanisms. Moreover, note that the optimal mechanism characterized above by altering the inspection probabilities of an optimal favored-agent mechanism is Pareto robustly dominated by the same optimal favored-agent mechanism. Thus, robust optimality alone is not a sufficient differentiator to distinguish between desirable and undesirable mechanisms. This insight prompts us to seek Pareto robustly

optimal mechanisms for problem (2.3). Next theorem shows that a favored-agent mechanism, proven to be optimal in Theorem 1, is also Pareto robustly optimal.

Theorem 2. *Any favored-agent mechanism of type (i) with favored agent $i^* \in \arg \max_{i \in \mathcal{I}} \underline{t}_i$ and threshold value $\nu^* = \max_{i \in \mathcal{I}} \underline{t}_i$ is Pareto robustly optimal in problem (2.3).*

We sketch the proof idea in the special case when there are only two agents. To convey the key ideas without tedious case distinctions, we assume that $\underline{t}_1 > \underline{t}_2$ so that $\arg \max_{i \in \mathcal{I}} \underline{t}_i = \{1\}$ is a singleton, and we assume that $\bar{t}_2 > c_2 + \underline{t}_1$ and $\bar{t}_1 > c_2 + \underline{t}_1$. We will use the following partition of the type space \mathcal{T} .

$$\begin{aligned}\mathcal{T}_I &= \{\mathbf{t} \in \mathcal{T} \mid t_2 - c_2 \leq \underline{t}_1 \text{ and } t_2 < \underline{t}_1\} \\ \mathcal{T}_{II} &= \{\mathbf{t} \in \mathcal{T} \mid t_2 - c_2 \leq \underline{t}_1 \text{ and } t_2 \geq \underline{t}_1\} \\ \mathcal{T}_{III} &= \{\mathbf{t} \in \mathcal{T} \mid t_2 - c_2 > \underline{t}_1 \text{ and } t_2 - c_2 > \underline{t}_1\} \\ \mathcal{T}_{IV} &= \{\mathbf{t} \in \mathcal{T} \mid t_2 - c_2 > \underline{t}_1 \text{ and } t_2 - c_2 \leq \underline{t}_1\}\end{aligned}$$

The sets \mathcal{T}_I , \mathcal{T}_{II} , \mathcal{T}_{III} and \mathcal{T}_{IV} are visualized in Figure 2.1. Note that all of them are nonempty thanks to our standing assumptions about \underline{t}_1 , \underline{t}_2 and c_2 . We emphasize, however, that all simplifying assumptions as well as the restriction to two agents are relaxed in the formal proof of Theorem 2.

In the following we denote by $(\mathbf{p}^*, \mathbf{q}^*)$ the favored-agent mechanism of type (i) with favored agent 1 and threshold value $\nu^* = \underline{t}_1$, and we will prove that this mechanism is Pareto robustly optimal in problem (2.3). To this end, assume for the sake of contradiction that there exists another mechanism (\mathbf{p}, \mathbf{q}) feasible in (2.3) that Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$. Thus, we have

$$\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) \geq \sum_{i \in \mathcal{I}} (p_i^*(\mathbf{t})t_i - q_i^*(\mathbf{t})c_i) \quad \forall \mathbf{t} \in \mathcal{T}, \quad (2.4)$$

where the inequality is strict for at least one $\mathbf{t} \in \mathcal{T}$. The right hand side of (2.4) represents

the principal's payoff in scenario \mathbf{t} under $(\mathbf{p}^*, \mathbf{q}^*)$. By the definition of a type (i) favored-agent mechanism, this payoff amounts to t_1 when $t_2 - c_2 \leq \underline{t}_1$ (i.e., when $\mathbf{t} \in \mathcal{T}_I \cup \mathcal{T}_{II}$) and to $\max_{i \in \mathcal{I}} t_i - c_i$ when $t_2 - c_2 > \underline{t}_1$ (i.e., when $\mathbf{t} \in \mathcal{T}_{III} \cup \mathcal{T}_{IV}$). We will show that if (2.4) holds, then (\mathbf{p}, \mathbf{q}) must generate the same payoff as $(\mathbf{p}^*, \mathbf{q}^*)$ under every type profile $\mathbf{t} \in \mathcal{T}$. In other words, (\mathbf{p}, \mathbf{q}) cannot generate a strictly higher payoff than $(\mathbf{p}^*, \mathbf{q}^*)$ under any type profile, which contradicts our assumption that (\mathbf{p}, \mathbf{q}) Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$.

We now investigate the subsets $\mathcal{T}_I, \mathcal{T}_{II}, \mathcal{T}_{III}$ and \mathcal{T}_{IV} of the type space one by one. Consider first a type profile $\mathbf{t} \in \mathcal{T}_I$. For inequality (2.4) to hold in this scenario, the principal must earn at least t_1 under the mechanism (\mathbf{p}, \mathbf{q}) . As $t_2 < t_1$, $c_i > 0$ and (\mathbf{p}, \mathbf{q}) satisfies the (FC) constraints $\sum_{i \in \mathcal{I}} p_i(\mathbf{t}) \leq 1$ and $q_i(\mathbf{t}) \geq 0$, this is only possible if $p_1(\mathbf{t}) = 1$ and $q_1(\mathbf{t}) = 0$. Thus, the allocation probabilities of the mechanisms (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ coincide on \mathcal{T}_I .

Consider now any $\mathbf{t} \in \mathcal{T}_{II}$. For inequality (2.4) to hold in scenario \mathbf{t} , the principal must earn at least t_1 under the mechanism (\mathbf{p}, \mathbf{q}) . Incentive compatibility ensures that $p_1(\mathbf{t}) \geq p_1(\bar{t}_1, t_2) - q_1(\bar{t}_1, t_2) = 1$, where the equality holds because $(\bar{t}_1, t_2) \in \mathcal{T}_I$ thanks to the assumption $\bar{t}_1 > c_2 + \underline{t}_1$ and because we know from before that (\mathbf{p}, \mathbf{q}) allocates the good to agent 1 without inspection in \mathcal{T}_I . Thus, the mechanism (\mathbf{p}, \mathbf{q}) can only earn t_1 in scenario \mathbf{t} if $p_1(\mathbf{t}) = 1$ and $q_1(\mathbf{t}) = 0$. In summary, the allocation probabilities of (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ must again coincide on \mathcal{T}_{II} .

Next, consider any $\mathbf{t} \in \mathcal{T}_{III}$. Incentive compatibility ensures that $p_2(\mathbf{t}) - q_2(\mathbf{t}) \leq$




Figure 2.1: Partition of the type space \mathcal{T} , (Support-Only).

$p_2(t_1, t_2) = 0$, where the equality holds because $(t_1, t_2) \in \mathcal{T}_I \cup \mathcal{T}_{II}$ and because we know from before that (\mathbf{p}, \mathbf{q}) allocates the good to agent 1 without inspection throughout $\mathcal{T}_I \cup \mathcal{T}_{II}$. As the allocation probabilities are non-negative and satisfy the (FC) condition $p_2(\mathbf{t}) \geq q_2(\mathbf{t})$, we may conclude that $p_2(\mathbf{t}) = q_2(\mathbf{t})$. Thus, the type of agent 2 is inspected if he wins the good in scenario \mathbf{t} . As $t_2 - c_2 > t_1 > t_1 - c_1$ for all $\mathbf{t} \in \mathcal{T}_{III}$, the inequality (2.4) implies that the principal must earn at least $t_2 - c_2$ under the mechanism (\mathbf{p}, \mathbf{q}) in scenario \mathbf{t} . This is only possible if $p_2(\mathbf{t}) = q_2(\mathbf{t}) = 1$. In summary, the allocation probabilities of (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ must therefore also coincide in \mathcal{T}_{III} .

Finally, consider any $\mathbf{t} \in \mathcal{T}_{IV}$. Incentive compatibility ensures that $0 = p_1(t_1, t_2) \geq p_1(\mathbf{t}) - q_1(\mathbf{t})$, where the equality holds because $(t_1, t_2) \in \mathcal{T}_{III}$ and because we know from before that (\mathbf{p}, \mathbf{q}) allocates the good to agent 2 in \mathcal{T}_{III} . Incentive compatibility also ensures that $0 = p_2(t_1, t_2) \geq p_2(\mathbf{t}) - q_2(\mathbf{t})$, where the equality holds because $(t_1, t_2) \in \mathcal{T}_I \cup \mathcal{T}_{II}$ and because we know from before that (\mathbf{p}, \mathbf{q}) allocates the good to agent 1 in $\mathcal{T}_I \cup \mathcal{T}_{II}$. We may thus conclude that $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \mathcal{I} = \{1, 2\}$. For the inequality (2.4) to hold in scenario \mathbf{t} , the principal must earn at least $\max_{i \in \mathcal{I}} t_i - c_i$ under the mechanism (\mathbf{p}, \mathbf{q}) . As $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \mathcal{I}$, this is only possible if (\mathbf{p}, \mathbf{q}) allocates the good to an agent $i' \in \arg \max_{i \in \mathcal{I}} t_i - c_i$ and inspects this agent. Thus, the principal's payoff under (\mathbf{p}, \mathbf{q}) matches her payoff under $(\mathbf{p}^*, \mathbf{q}^*)$ in region \mathcal{T}_{IV} .

The above reasoning shows that the principal's earnings coincide under (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ throughout the entire type space \mathcal{T} . Therefore, (\mathbf{p}, \mathbf{q}) cannot Pareto robustly dominate $(\mathbf{p}^*, \mathbf{q}^*)$, which in turn proves that $(\mathbf{p}^*, \mathbf{q}^*)$ is Pareto robustly optimal in problem (2.3).

Proof of Theorem 2. Throughout the proof we use the following partition of the type

space \mathcal{T} .

$$\mathcal{T}_I = \{\mathbf{t} \in \mathcal{T} \mid \max_{i \neq i^*} t_i - c_i \leq \underline{t}_{i^*} \text{ and } \max_{i \neq i^*} t_i < t_{i^*}\}$$

$$\mathcal{T}_{II} = \{\mathbf{t} \in \mathcal{T} \mid \max_{i \neq i^*} t_i - c_i \leq \underline{t}_{i^*} \text{ and } \max_{i \neq i^*} t_i \geq t_{i^*}\}$$

$$\mathcal{T}_{III} = \{\mathbf{t} \in \mathcal{T} \mid \max_{i \neq i^*} t_i - c_i > \underline{t}_{i^*} \text{ and } t_i - c_i \notin (\underline{t}_{i^*}, t_{i^*}] \forall i \neq i^*\}$$

$$\mathcal{T}_{IV} = \{\mathbf{t} \in \mathcal{T} \mid \max_{i \neq i^*} t_i - c_i > \underline{t}_{i^*} \text{ and } \exists i \neq i^* \text{ such that } t_i - c_i \in (\underline{t}_{i^*}, t_{i^*}]\}$$

Note that the set \mathcal{T}_I is nonempty and contains at least $(\bar{t}_{i^*}, \underline{t}_{-i^*})$ since $\max_{i \in \mathcal{I}} t_i = \underline{t}_{i^*} < \bar{t}_{i^*}$. However, the sets \mathcal{T}_{II} , \mathcal{T}_{III} and \mathcal{T}_{IV} can be empty if \underline{t}_{i^*} or c_i , $i \neq i^*$, are sufficiently large.

In the following, we denote by $(\mathbf{p}^*, \mathbf{q}^*)$ the favored-agent mechanism of type (i) with favored agent $i^* \in \arg \max_{i \in \mathcal{I}} \underline{t}_i$ and threshold value $\nu^* = \max_{i \in \mathcal{I}} \underline{t}_i$. By construction, we thus have $\nu^* = \underline{t}_{i^*}$. Assume now for the sake of contradiction that there exists another mechanism $(\mathbf{p}, \mathbf{q}) \in \mathcal{X}$ that Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$. Thus, the inequality (2.4) holds for all $\mathbf{t} \in \mathcal{T}$ and is strict for at least one $\mathbf{t} \in \mathcal{T}$. Note that the right-hand side of (2.4) represents the principal's payoff in scenario \mathbf{t} under $(\mathbf{p}^*, \mathbf{q}^*)$. By the definition of a type (i) favored-agent mechanism, this payoff amounts to t_{i^*} when $\max_{i \neq i^*} t_i - c_i \leq \underline{t}_{i^*}$ (i.e., when $\mathbf{t} \in \mathcal{T}_I \cup \mathcal{T}_{II}$) and to $\max_{i \in \mathcal{I}} t_i - c_i$ when $\max_{i \neq i^*} t_i - c_i > \underline{t}_{i^*}$ (i.e., when $\mathbf{t} \in \mathcal{T}_{III} \cup \mathcal{T}_{IV}$). We will show that if (2.4) holds, then (\mathbf{p}, \mathbf{q}) must generate the same payoff as $(\mathbf{p}^*, \mathbf{q}^*)$ under every type profile $\mathbf{t} \in \mathcal{T}$. In other words, (\mathbf{p}, \mathbf{q}) cannot generate a strictly higher payoff than $(\mathbf{p}^*, \mathbf{q}^*)$ under any type profile, which contradicts our assumption that (\mathbf{p}, \mathbf{q}) Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$. The remainder of the proof is divided into four steps, each of which investigates one of the subsets \mathcal{T}_I , \mathcal{T}_{II} , \mathcal{T}_{III} and \mathcal{T}_{IV} .

Step 1 (\mathcal{T}_I). Consider any type profile $\mathbf{t} \in \mathcal{T}_I$. For inequality (2.4) to hold in this scenario, the principal must earn at least t_{i^*} under mechanism (\mathbf{p}, \mathbf{q}) . We next show that this is only possible if $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$. To this end, assume for the sake of contradiction that either $p_{i^*}(\mathbf{t}) < 1$ or $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) > 0$. If $p_{i^*}(\mathbf{t}) < 1$, then the

principal's payoff under (\mathbf{p}, \mathbf{q}) satisfies

$$\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) \leq \sum_{i \in \mathcal{I}} p_i(\mathbf{t})t_i < t_{i^*} = \sum_{i \in \mathcal{I}} (p_i^*(\mathbf{t})t_i - q_i^*(\mathbf{t})c_i),$$

where the strict inequality holds because $\mathbf{t} \in \mathcal{T}_I$, which implies that $t_i < t_{i^*}$ for all $i \neq i^*$. Thus, inequality (2.4) is violated in scenario \mathbf{t} . If $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) > 0$, on the other hand, we have

$$\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) = p_{i^*}(\mathbf{t})t_{i^*} - q_{i^*}(\mathbf{t})c_{i^*} < t_{i^*} = \sum_{i \in \mathcal{I}} (p_i^*(\mathbf{t})t_i - q_i^*(\mathbf{t})c_i),$$

where the strict inequality holds because $q_{i^*}(\mathbf{t})$ and c_{i^*} are positive. Thus, inequality (2.4) is again violated in scenario \mathbf{t} . For inequality (2.4) to hold, we must therefore have $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$. Thus, the allocation probabilities of the mechanisms (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ coincide on \mathcal{T}_I .

Step 2 (\mathcal{T}_{II}). For inequality (2.4) to hold in any scenario $\mathbf{t} \in \mathcal{T}_{II}$, the principal must earn at least t_{i^*} under mechanism (\mathbf{p}, \mathbf{q}) . As in Step 1, we can show that this is only possible if $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$. To this end, we partition \mathcal{T}_{II} into the following subsets.

$$\mathcal{T}_{II_1} = \{\mathbf{t} \in \mathcal{T}_{II} \mid \max_{i \neq i^*} t_i < \bar{t}_{i^*}\}$$

$$\mathcal{T}_{II_2} = \{\mathbf{t} \in \mathcal{T}_{II} \mid \max_{i \neq i^*} t_i \geq \bar{t}_{i^*} \text{ and } t_{i^*} = \bar{t}_{i^*}\}$$

$$\mathcal{T}_{II_3} = \{\mathbf{t} \in \mathcal{T}_{II} \mid \max_{i \neq i^*} t_i \geq \bar{t}_{i^*} \text{ and } t_{i^*} < \bar{t}_{i^*}\}$$

Note that if $\max_{i \neq i^*} \bar{t}_i < \bar{t}_{i^*}$, then \mathcal{T}_{II} as well as its subsets \mathcal{T}_{II_1} , \mathcal{T}_{II_2} and \mathcal{T}_{II_3} are all empty. If $\max_{i \neq i^*} \bar{t}_i \geq \bar{t}_{i^*}$, on the other hand, then \mathcal{T}_{II} and its subset \mathcal{T}_{II_1} are nonempty. Indeed, \mathcal{T}_{II_1} contains the type profile \mathbf{t} defined through $t_i = \min\{\bar{t}_{i^*}, \bar{t}_i\}$ for all $i \in \mathcal{I}$. To see this, note that $\mathbf{t} \in \mathcal{T}$ by the construction of i^* . In addition, we have $\mathbf{t} \in \mathcal{T}_{II_1}$ thanks to the assumption $\max_{i \neq i^*} \bar{t}_i \geq \bar{t}_{i^*}$, which implies that $\max_{i \neq i^*} t_i = \bar{t}_{i^*}$. We now investigate the sets \mathcal{T}_{II_1} , \mathcal{T}_{II_2} and \mathcal{T}_{II_3} one by one.

Fix first any type profile $\mathbf{t} \in \mathcal{T}_{II_1}$. Incentive compatibility ensures that $p_{i^*}(\mathbf{t}) \geq p_{i^*}(\bar{t}_{i^*}, \mathbf{t}_{-i^*}) - q_{i^*}(\bar{t}_{i^*}, \mathbf{t}_{-i^*}) = 1$, where the equality holds because $(\bar{t}_{i^*}, \mathbf{t}_{-i^*}) \in \mathcal{T}_I$ and because we know from Step 1 that (\mathbf{p}, \mathbf{q}) allocates the good to agent i^* without inspection in \mathcal{T}_I . Consequently, the mechanism (\mathbf{p}, \mathbf{q}) can only earn t_{i^*} in scenario \mathbf{t} if $q_{i^*}(\mathbf{t}) = 0$. As $\mathbf{t} \in \mathcal{T}_{II_1}$ was chosen arbitrarily, the allocation probabilities (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ must therefore coincide throughout \mathcal{T}_{II_1} .

Next, we study the subset \mathcal{T}_{II_2} . To this end, define the set-valued function $\mathcal{I}(\mathbf{t}) = \{i \in \mathcal{I} \mid t_i \geq \bar{t}_{i^*}\}$ for $\mathbf{t} \in \mathcal{T}_{II_2}$. Note that $|\mathcal{I}(\mathbf{t})| \geq 2$ for all $\mathbf{t} \in \mathcal{T}_{II_2}$ thanks to the definition of \mathcal{T}_{II_2} , which implies that $i^* \in \mathcal{I}(\mathbf{t})$ and $\arg \max_{i \neq i^*} t_i \subseteq \mathcal{I}(\mathbf{t})$. We now prove by induction that the allocation probabilities (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ must coincide on $\mathcal{T}_{II_2}^n = \{\mathbf{t} \in \mathcal{T}_{II_2} \mid |\mathcal{I}(\mathbf{t})| = n\}$ for all $n \geq 2$.

As for the base step, set $n = 2$ and fix any type profile $\mathbf{t} \in \mathcal{T}_{II_2}^2$. Thus, there exists exactly one agent $i^\circ \neq i^*$ with type $t_{i^\circ} \geq \bar{t}_{i^*}$. Incentive compatibility dictates that $p_{i^\circ}(\mathbf{t}) - q_{i^\circ}(\mathbf{t}) \leq p_{i^\circ}(\underline{t}_{i^\circ}, \mathbf{t}_{-i^\circ}) = 0$, where the equality holds because $(\underline{t}_{i^\circ}, \mathbf{t}_{-i^\circ}) \in \mathcal{T}_I$ and because we know from Step 1 that (\mathbf{p}, \mathbf{q}) allocates the good to agent i^* without inspection in \mathcal{T}_I . We thus have $p_{i^\circ}(\mathbf{t}) = q_{i^\circ}(\mathbf{t})$. Inequality (2.4) further requires the mechanism (\mathbf{p}, \mathbf{q}) to earn at least \bar{t}_{i^*} in scenario $\mathbf{t} \in \mathcal{T}_{II_2}^2$. All of this is only possible if $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$ because $t_{i^\circ} - c_{i^\circ} \leq \underline{t}_{i^*} < \bar{t}_{i^*}$ and $t_i < \bar{t}_{i^*}$ for all $i \in \mathcal{I} \setminus \{i^\circ, i^*\}$.

As for the induction step, assume that $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$ for all $\mathbf{t} \in \mathcal{T}_{II_2}^n$ and for some $n \geq 2$, and fix an arbitrary type profile $\mathbf{t} \in \mathcal{T}_{II_2}^{n+1}$. Thus, there exist exactly n agents $i \neq i^*$ with types $t_i \geq \bar{t}_{i^*}$. For any such agent i , incentive compatibility dictates that $p_i(\mathbf{t}) - q_i(\mathbf{t}) \leq p_i(\underline{t}_i, \mathbf{t}_{-i}) = 0$, where the equality follows from the induction hypothesis and the observation that $(\underline{t}_i, \mathbf{t}_{-i}) \in \mathcal{T}_{II_2}^n$. We thus have $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \mathcal{I}(\mathbf{t}) \setminus \{i^*\}$. Inequality (2.4) further requires the mechanism (\mathbf{p}, \mathbf{q}) to earn at least \bar{t}_{i^*} in scenario $\mathbf{t} \in \mathcal{T}_{II_2}^{n+1}$. In analogy to the base step, all of this is only possible if $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$ because $t_i - c_i \leq \underline{t}_{i^*} < \bar{t}_{i^*}$ for all $i \in \mathcal{I}(\mathbf{t}) \setminus \{i^*\}$ and $t_i < \bar{t}_{i^*}$ for all $i \in \mathcal{I} \setminus \mathcal{I}(\mathbf{t})$. This observation completes the induction step. In summary, the allocation probabilities (\mathbf{p}, \mathbf{q})

and $(\mathbf{p}^*, \mathbf{q}^*)$ must therefore coincide throughout $\cup_{n \geq 2} \mathcal{T}_{II_2}^n = \mathcal{T}_{II_2}$.

Finally, fix any type profile $\mathbf{t} \in \mathcal{T}_{II_3}$. Incentive compatibility ensures that $p_{i^*}(\mathbf{t}) \geq p_{i^*}(\bar{t}_{i^*}, \mathbf{t}_{-i^*}) - q_{i^*}(\bar{t}_{i^*}, \mathbf{t}_{-i^*}) = 1$, where the equality holds because $(\bar{t}_{i^*}, \mathbf{t}_{-i^*}) \in \mathcal{T}_{II_2}$ and because we know from the above induction argument that (\mathbf{p}, \mathbf{q}) allocates the good to agent i^* without inspection in \mathcal{T}_{II_2} . Hence, the mechanism (\mathbf{p}, \mathbf{q}) can only earn t_{i^*} in scenario \mathbf{t} if $q_{i^*}(\mathbf{t}) = 0$. As $\mathbf{t} \in \mathcal{T}_{II_3}$ was chosen arbitrarily, the allocation probabilities (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ must therefore coincide throughout \mathcal{T}_{II_3} .

Step 3 (\mathcal{T}_{III}). In this part of the proof we will demonstrate that

$$\sum_{i \in \arg \max_{i' \in \mathcal{I}} t_{i'} - c_{i'}} p_i(\mathbf{t}) = 1 \quad \text{and} \quad p_i(\mathbf{t}) = q_i(\mathbf{t}) \quad \forall i \in \arg \max_{i' \in \mathcal{I}} t_{i'} - c_{i'}, \quad (2.5)$$

for every fixed $\mathbf{t} \in \mathcal{T}_{III}$. To prove (2.5), define the set-valued function $\mathcal{I}(\mathbf{t}) = \{i \in \mathcal{I} \mid t_i > t_{i^*}\}$ for $\mathbf{t} \in \mathcal{T}_{III}$. Note that $|\mathcal{I}(\mathbf{t})| \geq 1$ for all $\mathbf{t} \in \mathcal{T}_{III}$ thanks to the definition of \mathcal{T}_{III} , which ensures that there exists at least one agent $i \in \mathcal{I}$ with $t_i - c_i > t_{i^*}$. We will now use induction to prove that (2.5) holds for all type profiles in $\mathcal{T}_{III}^n = \{\mathbf{t} \in \mathcal{T}_{III} \mid |\mathcal{I}(\mathbf{t})| = n\}$ for all $n \geq 1$.

As for the base step, set $n = 1$ and fix any type profile $\mathbf{t} \in \mathcal{T}_{III}^1$. Thus, there exists exactly one agent $i^\circ \neq i^*$ with $t_{i^\circ} > t_{i^*}$. Incentive compatibility ensures that $p_{i^\circ}(\mathbf{t}) - q_{i^\circ}(\mathbf{t}) \leq p_{i^\circ}(t_{i^\circ}, \mathbf{t}_{-i^\circ}) = 0$, where the equality holds because $(t_{i^\circ}, \mathbf{t}_{-i^\circ}) \in \mathcal{T}_I \cup \mathcal{T}_{II}$. We thus have $p_{i^\circ}(\mathbf{t}) = q_{i^\circ}(\mathbf{t})$. If $p_{i^\circ}(\mathbf{t}) < 1$, then

$$\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) \leq p_{i^\circ}(\mathbf{t})(t_{i^\circ} - c_{i^\circ}) + \sum_{i \neq i^\circ} p_i(\mathbf{t})t_i < \max_{i \in \mathcal{I}} t_i - c_i,$$

where the first inequality holds because $p_{i^\circ}(\mathbf{t}) = q_{i^\circ}(\mathbf{t})$ and $c_i > 0$ for all $i \neq i^\circ$. The second inequality follows from the assumption that $p_{i^\circ}(\mathbf{t}) < 1$ as well as the definition of \mathcal{T}_{III} and the construction of i° , which imply that $t_{i^\circ} - c_{i^\circ} = \max_{i \in \mathcal{I}} t_i - c_i > t_{i^*}$ and $t_{i^*} \geq t_i$ for all $i \neq i^\circ$. This shows that (\mathbf{p}, \mathbf{q}) earns strictly less than $(\mathbf{p}^*, \mathbf{q}^*)$ in scenario \mathbf{t} ,

which contradicts inequality (2.4). Hence, our assumption must have been wrong, and $p_{i^*}(\mathbf{t})$ must equal 1. We have thus established (2.5) in scenario \mathbf{t} .

As for the induction step, assume that (2.5) holds throughout \mathcal{T}_{III}^n for some $n \geq 1$, and fix any type profile $\mathbf{t} \in \mathcal{T}_{III}^{n+1}$. Thus, there exist exactly $n + 1$ agents $i \neq i^*$ with types $t_i > t_{i^*}$. For any agent $i \in \mathcal{I}(\mathbf{t})$ incentive compatibility dictates that $p_i(\mathbf{t}) - q_i(\mathbf{t}) \leq p_i(\underline{t}_i, \mathbf{t}_{-i}) = 0$, where the equality holds because $(\underline{t}_i, \mathbf{t}_{-i}) \in \mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III}^n$. Indeed, if $(\underline{t}_i, \mathbf{t}_{-i}) \in \mathcal{T}_I \cup \mathcal{T}_{II}$, then the equality follows from the results of Steps 1 and 2, and if $(\underline{t}_i, \mathbf{t}_{-i}) \in \mathcal{T}_{III}^n$, then the equality follows from the induction hypothesis. We thus have $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \mathcal{I}(\mathbf{t})$ and, by the definition of \mathcal{T}_{III} , in particular for all $i \in \arg \max_{j \in \mathcal{I}} t_j - c_j$. In addition, if the summation of $p_i(\mathbf{t})$ over all $i \in \arg \max_{i' \in \mathcal{I}} t_{i'} - c_{i'}$ is strictly smaller than 1, then

$$\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) \leq \sum_{i \in \mathcal{I}(\mathbf{t})} p_i(\mathbf{t})(t_i - c_i) + \sum_{i \notin \mathcal{I}(\mathbf{t})} p_i(\mathbf{t})t_i < \max_{i \in \mathcal{I}} t_i - c_i$$

where the first inequality holds because $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \mathcal{I}(\mathbf{t})$ and $c_i > 0$ for all $i \notin \mathcal{I}(\mathbf{t})$. The strict inequality holds because $\sum_{i \in \mathcal{I}} p_i(\mathbf{t}) \leq 1$ and $\max_{j \in \mathcal{I}} t_j - c_j > t_{i^*} \geq t_i$ for all $i \notin \mathcal{I}(\mathbf{t})$ by the definition of \mathcal{T}_{III} and because we assumed that the summation of $p_i(\mathbf{t})$ over $i \in \arg \max_{i' \in \mathcal{I}} t_{i'} - c_{i'}$ is strictly smaller than 1. This reasoning shows that (\mathbf{p}, \mathbf{q}) earns strictly less than $(\mathbf{p}^*, \mathbf{q}^*)$ in scenario \mathbf{t} , which contradicts inequality (2.4). Hence, our assumption must be false and the summation of $p_i(\mathbf{t})$ over $i \in \arg \max_{i' \in \mathcal{I}} t_{i'} - c_{i'}$ equals 1. We have thus established (2.5) in scenario \mathbf{t} . As $\mathbf{t} \in \mathcal{T}_{III}^{n+1}$ was chosen arbitrarily, we may conclude that (2.5) holds throughout \mathcal{T}_{III}^{n+1} . This observation completes the induction step. In summary, the revenues generated by the mechanisms (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ must therefore coincide throughout $\cup_{n \geq 1} \mathcal{T}_{III}^n = \mathcal{T}_{III}$.

Step 4 (\mathcal{T}_{IV}). In analogy to Step 3, we will show that (2.5) holds for every fixed $\mathbf{t} \in \mathcal{T}_{IV}$. This immediately implies that (\mathbf{p}, \mathbf{q}) generates the same payoff as $(\mathbf{p}^*, \mathbf{q}^*)$ throughout \mathcal{T}_{IV} . To prove (2.5), define the set-valued function $\mathcal{I}(\mathbf{t}) = \{i \in \mathcal{I} \mid t_i > \underline{t}_{i^*}\}$ for $\mathbf{t} \in \mathcal{T}_{IV}$. Note that $|\mathcal{I}(\mathbf{t})| \geq 2$ for all $\mathbf{t} \in \mathcal{T}_{IV}$ thanks to the definition of \mathcal{T}_{IV} , which implies that

$i^* \in \mathcal{I}(\mathbf{t})$ and $\arg \max_{i \neq i^*} t_i - c_i \subseteq \mathcal{I}(\mathbf{t})$. To see that $i^* \in \mathcal{I}(\mathbf{t})$, note that if $i^* \notin \mathcal{I}(\mathbf{t})$ for some $\mathbf{t} \in \mathcal{T}_{IV}$, then $t_{i^*} = \underline{t}_{i^*}$, and there can be no $i \neq i^*$ with $t_i - c_i \in (\underline{t}_{i^*}, t_{i^*}] = \emptyset$, which contradicts the assumption that $\mathbf{t} \in \mathcal{T}_{IV}$. We will now use induction to prove that (2.5) holds for all type profiles in $\mathcal{T}_{IV}^n = \{\mathbf{t} \in \mathcal{T}_{IV} \mid |\mathcal{I}(\mathbf{t})| = n\}$ for all $n \geq 2$.

As for the base step, set $n = 2$ and fix an arbitrary type profile $\mathbf{t} \in \mathcal{T}_{IV}^2$. Thus, there exists exactly one agent $i^\circ \neq i^*$ with $t_{i^\circ} > \underline{t}_{i^*}$. Incentive compatibility for agent i^* ensures that $p_{i^*}(\mathbf{t}) - q_{i^*}(\mathbf{t}) \leq p_{i^*}(\underline{t}_{i^*}, \mathbf{t}_{-i^*}) = 0$, where the equality follows from (2.5) and the observation that $(\underline{t}_{i^*}, \mathbf{t}_{-i^*}) \in \mathcal{T}_{III}$. Thus, we have $p_{i^*}(\mathbf{t}) = q_{i^*}(\mathbf{t})$. Incentive compatibility for agent i° further dictates that $p_{i^\circ}(\mathbf{t}) - q_{i^\circ}(\mathbf{t}) \leq p_{i^\circ}(\underline{t}_{i^\circ}, \mathbf{t}_{-i^\circ}) = 0$, where the equality holds because $(\underline{t}_{i^\circ}, \mathbf{t}_{-i^\circ}) \in \mathcal{T}_I \cup \mathcal{T}_{II}$. Indeed, recall that the allocation probabilities of (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ match and that the good is allocated to agent i^* on $\mathcal{T}_I \cup \mathcal{T}_{II}$. Thus, we have $p_{i^\circ}(\mathbf{t}) = q_{i^\circ}(\mathbf{t})$. This reasoning shows that $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \mathcal{I}(\mathbf{t})$. Assume now that the summation of $p_i(\mathbf{t})$ over all $i \in \arg \max_{i' \in \mathcal{I}} t_{i'} - c_{i'}$ is strictly smaller than 1. Then, we have

$$\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) \leq \sum_{i \in \mathcal{I}(\mathbf{t})} p_i(\mathbf{t})(t_i - c_i) + \sum_{i \notin \mathcal{I}(\mathbf{t})} p_i(\mathbf{t})t_i < \max_{i \in \mathcal{I}} t_i - c_i,$$

where the first inequality holds because $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \mathcal{I}(\mathbf{t})$ and $c_i > 0$ for all $i \notin \mathcal{I}(\mathbf{t})$. The strict inequality holds because $\sum_{i \in \mathcal{I}} p_i(\mathbf{t}) \leq 1$, $\max_{j \in \mathcal{I}} t_j - c_j > \underline{t}_{i^*} \geq t_i$ for all $i \notin \mathcal{I}(\mathbf{t})$ by the definition of \mathcal{T}_{IV} and because we assumed that the summation of $p_i(\mathbf{t})$ over $i \in \arg \max_{i' \in \mathcal{I}} t_{i'} - c_{i'}$ is strictly smaller than 1. Hence, (\mathbf{p}, \mathbf{q}) earns strictly less than $(\mathbf{p}^*, \mathbf{q}^*)$ in scenario \mathbf{t} , which contradicts inequality (2.4). This implies that our assumption was false and that the summation of $p_i(\mathbf{t})$ over all $i \in \arg \max_{i' \in \mathcal{I}} t_{i'} - c_{i'}$ must be equal to 1. We have thus established (2.5) in scenario \mathbf{t} . As $\mathbf{t} \in \mathcal{T}_{IV}^2$ was chosen arbitrarily, (2.5) holds throughout \mathcal{T}_{IV}^2 .

As for the induction step, assume that (2.5) holds throughout \mathcal{T}_{IV}^n for some $n \geq 2$, and fix an arbitrary type profile $\mathbf{t} \in \mathcal{T}_{IV}^{n+1}$. Thus, there exist exactly n agents $i \neq i^*$ with types $t_i > \underline{t}_{i^*}$. Using the exact same reasoning as in the base step, we can prove

that $p_{i^*}(\mathbf{t}) = q_{i^*}(\mathbf{t})$. In addition, for any agent $i \in \mathcal{I}(\mathbf{t}) \setminus \{i^*\}$ incentive compatibility dictates that $p_i(\mathbf{t}) - q_i(\mathbf{t}) \leq p_i(\underline{t}_i, \mathbf{t}_{-i}) = 0$, where the equality holds because $(\underline{t}_i, \mathbf{t}_{-i}) \in \mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III} \cup \mathcal{T}_{IV}^n$. Indeed, if $(\underline{t}_i, \mathbf{t}_{-i}) \in \mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III}$, then the equality follows from the results of Steps 1, 2 and 3, and if $(\underline{t}_i, \mathbf{t}_{-i}) \in \mathcal{T}_{IV}^n$, then the equality follows from the induction hypothesis. In summary, we have thus shown that $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \mathcal{I}(\mathbf{t})$. The first statement in (2.5) can be proved by repeating the corresponding arguments from the base step almost verbatim. Details are omitted for brevity. We have thus established (2.5) in an arbitrary scenario $\mathbf{t} \in \mathcal{T}_{IV}^{n+1}$. By induction, the revenues generated by the mechanisms (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ must therefore coincide throughout $\cup_{n \geq 2} \mathcal{T}_{IV}^n = \mathcal{T}_{IV}$. This observation completes the proof. \square

2.4 Markov Ambiguity Sets

Although simple and adequate for situations where there is no distributional information, support-only ambiguity sets may be perceived conservative in practice. Motivated by this fact, we next investigate the mechanism design problem (MDP) under the assumption that the ambiguity set is a Markov ambiguity set of the form

$$\mathcal{P} = \{\mathbb{P} \in \mathcal{P}_0(\mathcal{T}) \mid \mathbb{E}_{\mathbb{P}}[\tilde{t}_i] \in [\underline{\mu}_i, \bar{\mu}_i] \ \forall i \in \mathcal{I}\}, \quad (2.6)$$

where $\underline{\mu}_i$ and $\bar{\mu}_i$ denote lower and upper bounds on the expected type $\mathbb{E}_{\mathbb{P}}[\tilde{t}_i]$ of agent $i \in \mathcal{I}$, respectively. We assume without much loss of generality that $\underline{t}_i < \underline{\mu}_i < \bar{\mu}_i < \bar{t}_i$ for all $i \in \mathcal{I}$. Under Markov ambiguity sets, the principal has information about the agent's mean types in addition to the support information.

Recall that if the principal knew the agents' types ex ante, she could simply allocate the good to the agent with the highest type without inspection. Therefore, the optimal value z^* of problem (MDP) is upper bounded by the value $\inf_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}}[\max_{i \in \mathcal{I}} \tilde{t}_i]$. In the

next proposition, we formally establish this upper bound and show that if \mathcal{P} is a Markov ambiguity set of the form (2.6), this upper bound amounts to $\max_{i \in \mathcal{I}} \underline{\mu}_i$ and is moreover attained by an admissible mechanism. Thus, this upper bound coincides with z^* .

Proposition 2. *If \mathcal{P} is a Markov ambiguity set of the form (2.6), then $z^* = \max_{i \in \mathcal{I}} \underline{\mu}_i$.*

Proof. Relaxing the incentive compatibility constraints and the first inequality in (FC) yields

$$\begin{aligned} z^* &\leq \sup_{\mathbf{p}, \mathbf{q}} \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} (p_i(\tilde{\mathbf{t}}) \tilde{t}_i - q_i(\tilde{\mathbf{t}}) c_i) \right] \\ &\text{s.t. } p_i : \mathcal{T} \rightarrow [0, 1] \text{ and } q_i : \mathcal{T} \rightarrow [0, 1] \quad \forall i \in \mathcal{I}, \\ &\quad \sum_{i \in \mathcal{I}} p_i(\mathbf{t}) \leq 1 \quad \forall \mathbf{t} \in \mathcal{T}, \\ &= \sup_{\mathbf{p}} \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} p_i(\tilde{\mathbf{t}}) \tilde{t}_i \right] \\ &\text{s.t. } p_i : \mathcal{T} \rightarrow [0, 1] \quad \forall i \in \mathcal{I}, \quad \sum_{i \in \mathcal{I}} p_i(\mathbf{t}) \leq 1 \quad \forall \mathbf{t} \in \mathcal{T}, \end{aligned}$$

where the equality holds because it is optimal to set $q_i(\mathbf{t}) = 0$ for all $i \in \mathcal{I}$ and $\mathbf{t} \in \mathcal{T}$ in the relaxed problem. As $p_i \geq 0$ and $\sum_{i \in \mathcal{I}} p_i(\mathbf{t}) \leq 1$ for all $\mathbf{t} \in \mathcal{T}$, we moreover have

$$\sum_{i \in \mathcal{I}} p_i(\mathbf{t}) t_i \leq \max_{i \in \mathcal{I}} t_i \quad \forall \mathbf{t} \in \mathcal{T},$$

which imply that z^* is bounded above by $\inf_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}} [\max_{i \in \mathcal{I}} \tilde{t}_i]$. Now, select an arbitrary $i^* \in \arg \max_{i \in \mathcal{I}} \underline{\mu}_i$ and denote by $\delta_{\underline{\mu}}$ the Dirac point mass at $\underline{\mu}$. We have

$$\mathbb{E}_{\delta_{\underline{\mu}}} \left[\max_{i \in \mathcal{I}} \tilde{t}_i \right] \geq \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}} \left[\max_{i \in \mathcal{I}} \tilde{t}_i \right] \geq \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}} [\tilde{t}_{i^*}] = \max_{i \in \mathcal{I}} \underline{\mu}_i,$$

where the first inequality holds because $\delta_{\underline{\mu}} \in \mathcal{P}$, the second inequality holds because $\max_{i \in \mathcal{I}} t_i \geq t_{i^*}$ for any $\mathbf{t} \in \mathcal{T}$, and the equality follows from the selection of i^* and the definition of the Markov ambiguity set \mathcal{P} . As $\delta_{\underline{\mu}}$ is the Dirac point mass at $\underline{\mu}$, we also have $\mathbb{E}_{\delta_{\underline{\mu}}} [\max_{i \in \mathcal{I}} \tilde{t}_i] = \max_{i \in \mathcal{I}} \underline{\mu}_i$ that implies $\inf_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}} [\max_{i \in \mathcal{I}} \tilde{t}_i] = \max_{i \in \mathcal{I}} \underline{\mu}_i$. Therefore,

the optimal value z^* is bounded above by $\max_{i \in \mathcal{I}} \underline{\mu}_i$. However, this bound is attained by a mechanism that allocates the good to an agent $i^* \in \arg \max_{i \in \mathcal{I}} \underline{\mu}_i$ irrespective of $\mathbf{t} \in \mathcal{T}$ and never inspects anyone's type. Since this mechanism is feasible, the claim follows. \square

Proposition 2 shows that the principal can secure a worst-case expected payoff of $\max_{i \in \mathcal{I}} \underline{\mu}_i$ under Markov ambiguity sets. Hence, in comparison to support-only information, the additional information about the mean types of the agents increases the principal's optimal worst-case expected payoff from $\max_{i \in \mathcal{I}} t_i$ to $\max_{i \in \mathcal{I}} \underline{\mu}_i$. In the next theorem, we characterize a class of favored-agent mechanisms that attain the optimal value $z^* = \max_{i \in \mathcal{I}} \underline{\mu}_i$ of problem (MDP) under Markov ambiguity sets.

Theorem 3. *If \mathcal{P} is a Markov ambiguity set of the form (2.6), then any favored-agent mechanism with favored agent $i^* \in \arg \max_{i \in \mathcal{I}} \underline{\mu}_i$ and threshold value $\nu^* \geq \bar{t}_{i^*}$ is optimal in (MDP).*

Proof. Select an arbitrary favored-agent mechanism with $i^* \in \arg \max_{i \in \mathcal{I}} \underline{\mu}_i$ and $\nu^* \geq \bar{t}_{i^*}$. Recall first that this mechanism is feasible in (MDP). Next, we will show that this mechanism attains a worst-case payoff that is at least as large as $\max_{i \in \mathcal{I}} \underline{\mu}_i$, which implies via Proposition 2 that this mechanism is optimal in (MDP). To this end, fix an arbitrary type profile $\mathbf{t} \in \mathcal{T}$. If $\max_{i \in \mathcal{I}} t_i - c_i < \nu^*$, then condition (i) in Definition 3 implies that the principal's payoff amounts to t_{i^*} . If $\max_{i \in \mathcal{I}} t_i - c_i > \nu^*$, then condition (ii) in Definition 3 implies that the principal's payoff amounts to $\max_{i \in \mathcal{I}} t_i - c_i > \nu^* \geq t_{i^*}$, where the second inequality follows from the selection of ν^* . If $\max_{i \neq i^*} t_i - c_i = \nu^*$, then the allocation functions are defined either as in condition (i) or as in condition (ii) of Definition 3. Thus, the principal's payoff amounts either to t_{i^*} or to $\max_{i \in \mathcal{I}} t_i - c_i \geq \nu^* \geq t_{i^*}$, respectively. In summary, we have shown that the principal's payoff is bigger than or equal to t_{i^*} in all three cases. As the type profile \mathbf{t} was chosen arbitrarily, this implies that the principal's expected payoff under any distribution $\mathbb{P} \in \mathcal{P}$ is bounded below by $\mathbb{E}_{\mathbb{P}} [\bar{t}_{i^*}]$. By the definition of the Markov ambiguity set \mathcal{P} , the expectation $\mathbb{E}_{\mathbb{P}} [\bar{t}_{i^*}]$ cannot be lower than $z^* = \max_{i \in \mathcal{I}} \underline{\mu}_i$ for any $\mathbb{P} \in \mathcal{P}$. Therefore, the principal's worst-case expected payoff under the favored-agent

mechanism is bounded below by z^* . The favored-agent mechanism at hand is therefore optimal in (2.3) by virtue of Proposition 2. \square

In the remainder of this section, we seek Pareto robustly optimal mechanisms for problem (MDP) under Markov ambiguity sets. To this end, we first present a set of preliminary results. Even though some of the following results rely on the assumption that the set $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i$ is a singleton, *i.e.*, there is a single candidate for the optimal favored agent, the Pareto robust optimality result of this section will not depend on this assumption.

Lemma 2. *If \mathcal{P} is a Markov ambiguity set of the form (2.6) and $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{i^*\}$ is a singleton, then, for any type profile $\mathbf{t} \in \mathcal{T}$, there exist a scenario $\hat{\mathbf{t}} \in \mathcal{T}$, where $\max_{i \neq i^*} \hat{t}_i < \hat{t}_{i^*}$, and a discrete distribution $\mathbb{P} \in \mathcal{P}$ that satisfy the following properties: (i) $\mathbb{E}_{\mathbb{P}}[\tilde{t}_i] = \underline{\mu}_i, \forall i \in \mathcal{I}$, (ii) $\mathbb{P}(\tilde{\mathbf{t}} \in \{\mathbf{t}, \hat{\mathbf{t}}\}) = 1$, (iii) $\mathbb{P}(\tilde{\mathbf{t}} = \mathbf{t}) > 0$.*

Proof. For any $\mathbf{t} \in \mathcal{T}$, we will show that there exists a scenario $\hat{\mathbf{t}} \in \mathcal{T}$ that satisfies $\max_{i \neq i^*} \hat{t}_i < \hat{t}_{i^*}$ and $\alpha \mathbf{t} + (1 - \alpha) \hat{\mathbf{t}} = \underline{\boldsymbol{\mu}}$ for some $\alpha \in (0, 1]$. This implies that the discrete distribution $\mathbb{P} = \alpha \delta_{\mathbf{t}} + (1 - \alpha) \delta_{\hat{\mathbf{t}}}$ belongs to the Markov ambiguity set \mathcal{P} and moreover satisfies the properties (i)–(iii).

To this end, consider any $\mathbf{t} \in \mathcal{T}$. If $\mathbf{t} = \underline{\boldsymbol{\mu}}$, set $\hat{\mathbf{t}} = \mathbf{t} = \underline{\boldsymbol{\mu}}$. As $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{i^*\}$ is a singleton, scenario $\hat{\mathbf{t}}$ satisfies $\max_{i \neq i^*} \hat{t}_i < \hat{t}_{i^*}$. Moreover, note that $\alpha \mathbf{t} + (1 - \alpha) \hat{\mathbf{t}} = \underline{\boldsymbol{\mu}}$ for any $\alpha \in (0, 1]$. Similarly, for any $\alpha \in (0, 1]$, $\mathbb{P} = \alpha \delta_{\mathbf{t}} + (1 - \alpha) \delta_{\hat{\mathbf{t}}} = \delta_{\underline{\boldsymbol{\mu}}}$ is the Dirac point mass at $\underline{\boldsymbol{\mu}}$ and trivially satisfies the desired properties (i)–(iii).

If $\mathbf{t} \neq \underline{\boldsymbol{\mu}}$, define function $\hat{\mathbf{t}}(\alpha)$ through

$$\hat{\mathbf{t}}(\alpha) = \frac{1}{1 - \alpha} (\underline{\boldsymbol{\mu}} - \mathbf{t}) + \mathbf{t}.$$

Note that, for any $\alpha \in [0, 1)$, $\hat{\mathbf{t}}(\alpha)$ satisfies

$$\alpha \mathbf{t} + (1 - \alpha) \hat{\mathbf{t}}(\alpha) = \alpha \mathbf{t} + (1 - \alpha) \left(\frac{1}{1 - \alpha} (\underline{\boldsymbol{\mu}} - \mathbf{t}) + \mathbf{t} \right) = \underline{\boldsymbol{\mu}}.$$

Thus, for any $\alpha \in [0, 1)$, $\hat{\mathbf{t}} = \hat{\mathbf{t}}(\alpha)$ satisfies $\alpha \mathbf{t} + (1 - \alpha) \hat{\mathbf{t}} = \underline{\boldsymbol{\mu}}$. We will next show that there exists an $\alpha \in (0, 1)$ for which $\hat{\mathbf{t}} = \hat{\mathbf{t}}(\alpha)$ also satisfies $\max_{i \neq i^*} \hat{t}_i < \hat{t}_{i^*}$. To this end, first note that $\hat{\mathbf{t}}(\alpha)$ is a continuous function of $\alpha \in [0, 1)$ and $\hat{\mathbf{t}}(0) = \underline{\boldsymbol{\mu}}$. Thus, for any $\varepsilon > 0$, there exists $\alpha \in (0, 1)$ such that $\hat{\mathbf{t}}(\alpha) \in \prod_{i \in \mathcal{I}} [\underline{\mu}_i - \varepsilon, \underline{\mu}_i + \varepsilon]$. We next show that any $\varepsilon > 0$ that belongs to the set

$$L = (0, \min_{i \in \mathcal{I}} \underline{\mu}_i - \underline{t}_i) \cap (0, \min_{i \in \mathcal{I}} \bar{t}_i - \underline{\mu}_i) \cap (0, (\underline{\mu}_{i^*} - \max_{i \neq i^*} \underline{\mu}_i)/2)$$

ensures that $\prod_{i \in \mathcal{I}} [\underline{\mu}_i - \varepsilon, \underline{\mu}_i + \varepsilon] \subseteq \{\mathbf{t} \in \mathcal{T} \mid \max_{i \neq i^*} t_i < t_{i^*}\}$. Note that set L is non-empty because $\underline{t}_i < \underline{\mu}_i < \bar{\mu}_i < \bar{t}_i$ for all $i \in \mathcal{I}$ and $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{i^*\}$ is a singleton. Consider any $\varepsilon \in L$. As $\varepsilon < \min_{i \in \mathcal{I}} \underline{\mu}_i - \underline{t}_i$, any $\mathbf{t} \in \prod_{i \in \mathcal{I}} [\underline{\mu}_i - \varepsilon, \underline{\mu}_i + \varepsilon]$ satisfies

$$t_i \geq \underline{\mu}_i - \varepsilon > \underline{\mu}_i - (\min_{j \in \mathcal{I}} \underline{\mu}_j - \underline{t}_j) \geq \underline{\mu}_i - (\underline{\mu}_i - \underline{t}_i) = \underline{t}_i \quad \forall i \in \mathcal{I}.$$

Similarly, as $\varepsilon < \min_{i \in \mathcal{I}} \bar{t}_i - \underline{\mu}_i$, any $\mathbf{t} \in \prod_{i \in \mathcal{I}} [\underline{\mu}_i - \varepsilon, \underline{\mu}_i + \varepsilon]$ satisfies

$$t_i \leq \underline{\mu}_i + \varepsilon < \underline{\mu}_i + (\min_{j \in \mathcal{I}} \bar{t}_j - \underline{\mu}_j) \leq \underline{\mu}_i + \bar{t}_i - \underline{\mu}_i = \bar{t}_i \quad \forall i \in \mathcal{I}.$$

Therefore, we have shown that $\prod_{i \in \mathcal{I}} [\underline{\mu}_i - \varepsilon, \underline{\mu}_i + \varepsilon] \subseteq \mathcal{T}$. Finally, any $\mathbf{t} \in \prod_{i \in \mathcal{I}} [\underline{\mu}_i - \varepsilon, \underline{\mu}_i + \varepsilon]$ satisfies

$$\begin{aligned} t_{i^*} &\geq \underline{\mu}_{i^*} - \varepsilon > \underline{\mu}_{i^*} - (\underline{\mu}_{i^*} - \max_{j \neq i^*} \underline{\mu}_j)/2 = (\underline{\mu}_{i^*} + \max_{j \neq i^*} \underline{\mu}_j)/2 \\ &= \max_{j \neq i^*} \underline{\mu}_j + (\underline{\mu}_{i^*} - \max_{j \neq i^*} \underline{\mu}_j)/2 > \max_{j \neq i^*} \underline{\mu}_j + \varepsilon \geq \underline{\mu}_i + \varepsilon \geq t_i \quad \forall i \neq i^*, \end{aligned}$$

where the second and third inequalities follow from $\varepsilon < (\underline{\mu}_{i^*} - \max_{i \neq i^*} \underline{\mu}_i)/2$. Thus, we have shown that $\prod_{i \in \mathcal{I}} [\underline{\mu}_i - \varepsilon, \underline{\mu}_i + \varepsilon] \subseteq \{\mathbf{t} \in \mathcal{T} \mid \max_{i \neq i^*} t_i < t_{i^*}\}$ for any $\varepsilon \in L$. As for any $\varepsilon \in L$ there exists $\alpha \in (0, 1)$ such that $\hat{\mathbf{t}}(\alpha) \in \prod_{i \in \mathcal{I}} [\underline{\mu}_i - \varepsilon, \underline{\mu}_i + \varepsilon]$, the claim follows. \square

The next proposition formalizes a necessary and sufficient optimality condition using

Lemma 2.

Proposition 3. *If \mathcal{P} is a Markov ambiguity set of the form (2.6) and the set $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{i^*\}$ is a singleton, then a mechanism $(\mathbf{p}, \mathbf{q}) \in \mathcal{X}$ is optimal in (MDP) if and only if*

$$\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) \geq t_{i^*} \quad \forall \mathbf{t} \in \mathcal{T}. \quad (2.7)$$

Proof. Consider an arbitrary mechanism $(\mathbf{p}, \mathbf{q}) \in \mathcal{X}$. If (\mathbf{p}, \mathbf{q}) satisfies (2.7), then the principal's expected payoff $\mathbb{E}_{\mathbb{P}} [\sum_{i \in \mathcal{I}} (p_i(\tilde{\mathbf{t}})\tilde{t}_i - q_i(\tilde{\mathbf{t}})c_i)]$ under any distribution $\mathbb{P} \in \mathcal{P}$ is at least $\mathbb{E}_{\mathbb{P}} [\tilde{t}_{i^*}] \geq \max_{i \in \mathcal{I}} \underline{\mu}_i$, where the inequality follows from the definition of the Markov ambiguity set \mathcal{P} . By virtue of Proposition 2, this mechanism is therefore optimal (MDP). We thus have shown that if (\mathbf{p}, \mathbf{q}) satisfies (2.7), then it is optimal in (MDP).

We next show that if (\mathbf{p}, \mathbf{q}) is optimal in (MDP), then it must satisfy (2.7). To this end, assume for the sake of contradiction that (\mathbf{p}, \mathbf{q}) is optimal and $\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) < t_{i^*}$ for some $\mathbf{t} \in \mathcal{T}$. Consider an arbitrary $\mathbf{t} \in \mathcal{T}$ for which inequality (2.7) fails. By Lemma 2, we know that there exist a scenario $\hat{\mathbf{t}} \in \mathcal{T}$, where $\max_{i \neq i^*} \hat{t}_i < \hat{t}_{i^*}$, and a discrete distribution $\mathbb{P} \in \mathcal{P}$ that satisfy the following properties: (i) $\mathbb{E}_{\mathbb{P}}[\tilde{t}_i] = \underline{\mu}_i, \forall i \in \mathcal{I}$, (ii) $\mathbb{P}(\tilde{\mathbf{t}} \in \{\mathbf{t}, \hat{\mathbf{t}}\}) = 1$, (iii) $\mathbb{P}(\tilde{\mathbf{t}} = \mathbf{t}) > 0$. The principal's payoff $\sum_{i \in \mathcal{I}} (p_i(\hat{\mathbf{t}})\hat{t}_i - q_i(\hat{\mathbf{t}})c_i)$ in scenario $\hat{\mathbf{t}}$ is bounded above by $\sum_{i \in \mathcal{I}} p_i(\hat{\mathbf{t}})\hat{t}_i \leq \hat{t}_{i^*}$, where the inequality holds because $\sum_{i \in \mathcal{I}} p_i(\hat{\mathbf{t}}) \leq 1$ and $\hat{t}_i \leq \hat{t}_{i^*}$ for all $i \in \mathcal{I}$. The principal's expected payoff under \mathbb{P} therefore satisfies

$$\begin{aligned} \mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} (p_i(\tilde{\mathbf{t}})\tilde{t}_i - q_i(\tilde{\mathbf{t}})c_i) \right] \\ = \mathbb{P}(\tilde{\mathbf{t}} = \mathbf{t}) \sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) + \mathbb{P}(\tilde{\mathbf{t}} = \hat{\mathbf{t}}) \sum_{i \in \mathcal{I}} (p_i(\hat{\mathbf{t}})\hat{t}_i - q_i(\hat{\mathbf{t}})c_i) \\ < \mathbb{P}(\tilde{\mathbf{t}} = \mathbf{t})t_{i^*} + \mathbb{P}(\tilde{\mathbf{t}} = \hat{\mathbf{t}})\hat{t}_{i^*} = \underline{\mu}_{i^*}, \end{aligned}$$

where the first equality follows from property (ii), the inequality holds because of property

(iii) and because we have assumed that $\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) < t_{i^*}$ and we have shown that $\sum_{i \in \mathcal{I}} (p_i(\hat{\mathbf{t}})\hat{t}_i - q_i(\hat{\mathbf{t}})c_i) \leq \hat{t}_{i^*}$, and the last equality follows from properties (i) and (ii). As the principal's expected payoff under \mathbb{P} is strictly smaller than $z^* = \underline{\mu}_{i^*}$, mechanism (\mathbf{p}, \mathbf{q}) cannot be optimal. The claim thus follows. \square

Proposition 3 reveals that the type t_{i^*} of agent i^* is an important reference point for optimality if $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{i^*\}$ is a singleton. A mechanism is optimal in (MDP) if and only if it earns a payoff of at least t_{i^*} under any type profile \mathbf{t} . Our next result shows that this optimality condition and incentive compatibility constraints uniquely determine the allocation probabilities of any optimal mechanism throughout a subset of all scenarios. In particular, an optimal mechanism in (MDP) should allocate the good to agent i^* without inspection if no other agent reports a value $t_i - c_i$ that exceeds the highest possible type \bar{t}_{i^*} of agent i^* .

Proposition 4. *If \mathcal{P} is a Markov ambiguity set of the form (2.6) and the set $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{i^*\}$ is a singleton, then any optimal mechanism (\mathbf{p}, \mathbf{q}) in (MDP) satisfies the following property. For any type profile $\mathbf{t} \in \mathcal{T}$ such that $\max_{i \neq i^*} t_i - c_i < \bar{t}_{i^*}$, $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$.*

We outline the proof idea in the special case when there are only two agents and when $\underline{\mu}_2 < \underline{\mu}_1$ so that $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{1\}$ is a singleton. We also assume that $\bar{t}_2 > c_2 + \bar{t}_1$ to prevent tedious case distinctions. Our arguments make use of the following partition of the type space \mathcal{T} .

$$\begin{aligned}
\mathcal{T}_I &= \{\mathbf{t} \in \mathcal{T} \mid t_2 < t_1\} \\
\mathcal{T}_{II} &= \{\mathbf{t} \in \mathcal{T} \mid t_2 \geq t_1 \text{ and } t_2 < \bar{t}_1\} \\
\mathcal{T}_{III} &= \{\mathbf{t} \in \mathcal{T} \mid t_2 \geq t_1, t_2 \geq \bar{t}_1 \text{ and } t_2 - c_2 < t_1\} \\
\mathcal{T}_{IV} &= \{\mathbf{t} \in \mathcal{T} \mid t_2 \geq t_1, t_2 \geq \bar{t}_1, t_2 - c_2 \geq t_1 \text{ and } t_2 - c_2 < \bar{t}_1\} \\
\mathcal{T}_V &= \{\mathbf{t} \in \mathcal{T} \mid t_2 \geq t_1, t_2 \geq \bar{t}_1, t_2 - c_2 \geq t_1 \text{ and } t_2 - c_2 \geq \bar{t}_1\}
\end{aligned} \tag{2.8}$$

Note that some of the conditions in set definitions above are redundant and given for ease of readability. Sets $\mathcal{T}_I - \mathcal{T}_V$ are illustrated in Figure 2.2. One can show that all of these sets are nonempty thanks to our standing assumptions about $\underline{\mu}_1, \underline{\mu}_2, \bar{t}_1, \bar{t}_2$ and c_2 . We emphasize, however, that all simplifying assumptions as well as the restriction to two agents are relaxed in the formal proof of Proposition 4.

In the following we use the optimality condition (2.7) that is given in Proposition 3, that is, any optimal mechanism must earn at least t_1 in any scenario $\mathbf{t} \in \mathcal{T}$. We will prove that, when agent 2 fails to report a type t_2 that is at least $c_2 + \bar{t}_1$, inequality (2.7) can be satisfied only if the good is allocated to agent 1 without inspection. Note that we have $t_2 < c_2 + \bar{t}_1$ under any scenario $\mathbf{t} \in \mathcal{T} \setminus \mathcal{T}_V$.

We now assume that a mechanism (\mathbf{p}, \mathbf{q}) is optimal and investigate the subsets $\mathcal{T}_I, \mathcal{T}_{II}, \mathcal{T}_{III}$ and \mathcal{T}_{IV} of the type space one by one. Consider first a type profile $\mathbf{t} \in \mathcal{T}_I$. As $t_2 < t_1$, $c_1 > 0$ and (\mathbf{p}, \mathbf{q}) satisfies the (FC) constraints $\sum_{i \in \mathcal{I}} p_i(\mathbf{t}) \leq 1$ and $q_1(\mathbf{t}) \geq 0$, the mechanism (\mathbf{p}, \mathbf{q}) can earn a payoff that is at least t_1 only if $p_1(\mathbf{t}) = 1$ and $q_1(\mathbf{t}) = 0$.

Consider now any $\mathbf{t} \in \mathcal{T}_{II}$. Incentive compatibility ensures that $p_1(\mathbf{t}) \geq p_1(\bar{t}_1, t_2) - q_1(\bar{t}_1, t_2) = 1$, where the equality holds because $(\bar{t}_1, t_2) \in \mathcal{T}_I$ and because we know from before that (\mathbf{p}, \mathbf{q}) allocates the good to agent 1 without inspection in \mathcal{T}_I . Thus, the mechanism (\mathbf{p}, \mathbf{q}) can earn a payoff that is at least t_1 in scenario \mathbf{t} only if $q_1(\mathbf{t}) = 0$. In summary, we must again have $p_1(\mathbf{t}) = 1$ and $q_1(\mathbf{t}) = 0$ so that (\mathbf{p}, \mathbf{q}) can satisfy (2.7).




Figure 2.2: Partition of the type space \mathcal{T} , (Markov).

Next, consider any scenario $\mathbf{t} \in \mathcal{T}_{III}$. Incentive compatibility ensures that $p_2(\mathbf{t}) - q_2(\mathbf{t}) \leq p_2(t_1, t_2) = 0$, where the equality holds because $(t_1, t_2) \in \mathcal{T}_I \cup \mathcal{T}_{II}$ and because we know from before that (\mathbf{p}, \mathbf{q}) allocates the good to agent 1 without inspection throughout $\mathcal{T}_I \cup \mathcal{T}_{II}$. As the allocation probabilities are non-negative and satisfy the (FC) condition $p_2(\mathbf{t}) \geq q_2(\mathbf{t})$, we may conclude that $p_2(\mathbf{t}) = q_2(\mathbf{t})$. Thus, the report of agent 2 is inspected if he wins the good in scenario \mathbf{t} . As $t_2 - c_2 < t_1$ for all $\mathbf{t} \in \mathcal{T}_{III}$, the mechanism (\mathbf{p}, \mathbf{q}) can earn a payoff that is at least t_1 only if $p_1(\mathbf{t}) = 1$ and $q_1(\mathbf{t}) = 0$.

Finally, consider any $\mathbf{t} \in \mathcal{T}_{IV}$. Incentive compatibility ensures that $p_1(\mathbf{t}) \geq p_1(\bar{t}_1, t_2) - q_1(\bar{t}_1, t_2) = 1$, where the equality holds because $(\bar{t}_1, t_2) \in \mathcal{T}_{III}$ and because we know from before that (\mathbf{p}, \mathbf{q}) allocates the good to agent 1 without inspection in \mathcal{T}_{III} . Thus, the mechanism (\mathbf{p}, \mathbf{q}) can earn a payoff that is at least t_1 in scenario \mathbf{t} only if $q_1(\mathbf{t}) = 0$. Hence, we must again have $p_1(\mathbf{t}) = 1$ and $q_1(\mathbf{t}) = 0$ so that (\mathbf{p}, \mathbf{q}) can satisfy (2.7).

The reasoning above shows that, under the assumption $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{1\}$, any optimal mechanism (\mathbf{p}, \mathbf{q}) should give the good to agent 1 without inspection in any scenario $\mathbf{t} \in \mathcal{T}$ that satisfies $t_2 - c_2 < \bar{t}_1$.

Proof of Proposition 4. Throughout the proof we use the following partition of the type space \mathcal{T} .

$$\begin{aligned}
\mathcal{T}_I &= \{\mathbf{t} \in \mathcal{T} \mid \max_{i \neq i^*} t_i < t_{i^*}\} \\
\mathcal{T}_{II} &= \{\mathbf{t} \in \mathcal{T} \mid \max_{i \neq i^*} t_i \geq t_{i^*} \text{ and } \max_{i \neq i^*} t_i < \bar{t}_{i^*}\} \\
\mathcal{T}_{III} &= \{\mathbf{t} \in \mathcal{T} \mid \max_{i \neq i^*} t_i \geq t_{i^*}, \max_{i \neq i^*} t_i \geq \bar{t}_{i^*} \text{ and } \max_{i \neq i^*} t_i - c_i < t_{i^*}\} \\
\mathcal{T}_{IV} &= \{\mathbf{t} \in \mathcal{T} \mid \max_{i \neq i^*} t_i \geq t_{i^*}, \max_{i \neq i^*} t_i \geq \bar{t}_{i^*}, \max_{i \neq i^*} t_i - c_i \geq t_{i^*} \text{ and } \max_{i \neq i^*} t_i - c_i < \bar{t}_{i^*}\} \\
\mathcal{T}_V &= \{\mathbf{t} \in \mathcal{T} \mid \max_{i \neq i^*} t_i \geq t_{i^*}, \max_{i \neq i^*} t_i \geq \bar{t}_{i^*}, \max_{i \neq i^*} t_i - c_i \geq t_{i^*} \text{ and } \max_{i \neq i^*} t_i - c_i \geq \bar{t}_{i^*}\}
\end{aligned} \tag{2.9}$$

Note again that some of the conditions in the definitions above are redundant and introduced for ease of readability. Note also that the set \mathcal{T}_I is nonempty and contains at

least $\underline{\mu} = (\mu_1, \dots, \mu_I)$ because $\arg \max_{i \in \mathcal{I}} \mu_i = \{i^*\}$ is a singleton. However, the sets $\mathcal{T}_{II}, \mathcal{T}_{III}, \mathcal{T}_{IV}$ and \mathcal{T}_V can be empty if \underline{t}_{i^*} or $c_i, i \neq i^*$, are sufficiently large.

In the following, we will use Proposition 3 that shows that any optimal mechanism should satisfy (2.7). In other words, any optimal mechanism should earn a payoff that is at least t_{i^*} in any scenario $\mathbf{t} \in \mathcal{T}$. To prove the claim, we will show that if a feasible mechanism (\mathbf{p}, \mathbf{q}) violates $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$ for some $\mathbf{t} \in \mathcal{T}$ such that $\max_{i \neq i^*} t_i - c_i < \bar{t}_{i^*}$, then it cannot satisfy (2.7). Consequently, mechanism (\mathbf{p}, \mathbf{q}) cannot be optimal. The remainder of the proof is divided into four steps, each of which investigates one of the subsets $\mathcal{T}_I, \mathcal{T}_{II}, \mathcal{T}_{III}$ and \mathcal{T}_{IV} . We have $\max_{i \neq i^*} t_i - c_i \geq \bar{t}_{i^*}$ for any $\mathbf{t} \in \mathcal{T}_V$, and for this reason we do not need to investigate this set.

Step 1 (\mathcal{T}_I). Assume for the sake of contradiction that a mechanism (\mathbf{p}, \mathbf{q}) is optimal in (MDP) and satisfy $p_{i^*}(\mathbf{t}) < 1$ or $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) > 0$ in some scenario $\mathbf{t} \in \mathcal{T}_I$. If $p_{i^*}(\mathbf{t}) < 1$, then the principal's payoff can be written as

$$\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) \leq \sum_{i \in \mathcal{I}} p_i(\mathbf{t})t_i < t_{i^*},$$

where the strict inequality holds because $\sum_{i \in \mathcal{I}} p_i(\mathbf{t}) \leq 1$ and $\mathbf{t} \in \mathcal{T}_I$, which implies that $t_i < t_{i^*}$ for all $i \neq i^*$. Thus, inequality (2.7) is violated in scenario \mathbf{t} . If $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) > 0$, on the other hand, we have

$$\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) = p_{i^*}(\mathbf{t})t_{i^*} - q_{i^*}(\mathbf{t})c_{i^*} < t_{i^*},$$

where the strict inequality holds because $q_{i^*}(\mathbf{t})$ and c_{i^*} are positive. Thus, inequality (2.7) is again violated in scenario \mathbf{t} . For inequality (2.7) to hold, we must therefore have $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$ for any $\mathbf{t} \in \mathcal{T}_I$.

Step 2 (\mathcal{T}_{II}). Consider any type profile $\mathbf{t} \in \mathcal{T}_{II}$. Incentive compatibility ensures that $p_{i^*}(\mathbf{t}) \geq p_{i^*}(\bar{t}_{i^*}, \mathbf{t}_{-i^*}) - q_{i^*}(\bar{t}_{i^*}, \mathbf{t}_{-i^*}) = 1$, where the equality holds because $(\bar{t}_{i^*}, \mathbf{t}_{-i^*}) \in \mathcal{T}_I$

and because we know from Step 1 that (\mathbf{p}, \mathbf{q}) allocates the good to agent i^* without inspection in \mathcal{T}_I . Consequently, a feasible mechanism (\mathbf{p}, \mathbf{q}) can earn at least t_{i^*} in scenario \mathbf{t} only if $q_{i^*}(\mathbf{t}) = 0$. As $\mathbf{t} \in \mathcal{T}_{II}$ was chosen arbitrarily, any optimal mechanism (\mathbf{p}, \mathbf{q}) should satisfy $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$ throughout \mathcal{T}_{II} .

Step 3 (\mathcal{T}_{III}). Define the set-valued function $\mathcal{I}(\mathbf{t}) = \{i \in \mathcal{I} \mid t_i \geq t_{i^*}\}$ for $\mathbf{t} \in \mathcal{T}_{III}$. Note that $|\mathcal{I}(\mathbf{t})| \geq 2$ for all $\mathbf{t} \in \mathcal{T}_{III}$ because $i^* \in \mathcal{I}(\mathbf{t})$ and because the definition of \mathcal{T}_{III} ensures that $\max_{i \neq i^*} t_i \geq t_{i^*}$. We now prove by induction that $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$ for all type profiles in $\mathcal{T}_{III}^n = \{\mathbf{t} \in \mathcal{T}_{III} \mid |\mathcal{I}(\mathbf{t})| = n\}$ for all $n \geq 2$.

As for the base step, set $n = 2$ and fix any $\mathbf{t} \in \mathcal{T}_{III}^2$. Thus, there exists exactly one agent $i^\circ \neq i^*$ with $t_{i^\circ} \geq t_{i^*}$. Incentive compatibility ensures that $p_{i^\circ}(\mathbf{t}) - q_{i^\circ}(\mathbf{t}) \leq p_{i^\circ}(\underline{t}_{i^\circ}, \mathbf{t}_{-i^\circ}) = 0$, where the equality holds because $(\underline{t}_{i^\circ}, \mathbf{t}_{-i^\circ}) \in \mathcal{T}_I \cup \mathcal{T}_{II}$ and because we know from Step 1 and 2 that (\mathbf{p}, \mathbf{q}) allocates the good to agent i^* without inspection in $\mathcal{T}_I \cup \mathcal{T}_{II}$. We thus have $p_{i^\circ}(\mathbf{t}) = q_{i^\circ}(\mathbf{t})$. As $t_{i^\circ} - c_{i^\circ} < t_{i^*}$ and $t_j < t_{i^*}$ for all $j \in \mathcal{I} \setminus \mathcal{I}(\mathbf{t})$, the mechanism (\mathbf{p}, \mathbf{q}) can satisfy the inequality (2.7) for $\mathbf{t} \in \mathcal{T}_{III}^2$ only if $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$.

As for the induction step, assume that $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$ for all $\mathbf{t} \in \mathcal{T}_{III}^n$ and for some $n \geq 2$, and fix an arbitrary type profile $\mathbf{t} \in \mathcal{T}_{III}^{n+1}$. Thus, there exist exactly n agents $i \neq i^*$ with types $t_i \geq t_{i^*}$. For any such agent i , incentive compatibility dictates that $p_i(\mathbf{t}) - q_i(\mathbf{t}) \leq p_i(\underline{t}_i, \mathbf{t}_{-i}) = 0$, where the equality holds because $(\underline{t}_i, \mathbf{t}_{-i}) \in \mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III}^n$. Indeed, if $(\underline{t}_i, \mathbf{t}_{-i}) \in \mathcal{T}_I \cup \mathcal{T}_{II}$, then the equality follows from the results of Steps 1 and 2, and if $(\underline{t}_i, \mathbf{t}_{-i}) \in \mathcal{T}_{III}^n$, then the equality follows from the induction hypothesis. We thus have $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \mathcal{I}(\mathbf{t}) \setminus \{i^*\}$. In analogy to the base step, a feasible mechanism (\mathbf{p}, \mathbf{q}) can satisfy the inequality (2.7) for $\mathbf{t} \in \mathcal{T}_{III}^{n+1}$ only if $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$ because $t_i - c_i < t_{i^*}$ for all $i \in \mathcal{I}(\mathbf{t}) \setminus \{i^*\}$, and $t_j < t_{i^*}$ for $j \in \mathcal{I} \setminus \mathcal{I}(\mathbf{t})$. This observation completes the induction step. In summary, the allocation probabilities of any optimal mechanism (\mathbf{p}, \mathbf{q}) should satisfy $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$ throughout $\cup_{n \geq 2} \mathcal{T}_{III}^n = \mathcal{T}_{III}$.

Step 4 (\mathcal{T}_{IV}). Fix now any arbitrary type profile $\mathbf{t} \in \mathcal{T}_{IV}$. Incentive compatibility ensures that $p_{i^*}(\mathbf{t}) \geq p_{i^*}(\bar{t}_{i^*}, \mathbf{t}_{-i^*}) - q_{i^*}(\bar{t}_{i^*}, \mathbf{t}_{-i^*}) = 1$, where the equality holds because $(\bar{t}_{i^*}, \mathbf{t}_{-i^*}) \in \mathcal{T}_{III}$ and because we know from Step 3 that any optimal mechanism (\mathbf{p}, \mathbf{q}) allocates the good to agent i^* without inspection in \mathcal{T}_{III} . Consequently, a feasible mechanism (\mathbf{p}, \mathbf{q}) can earn at least t_{i^*} in scenario \mathbf{t} only if $q_{i^*}(\mathbf{t}) = 0$. As $\mathbf{t} \in \mathcal{T}_{IV}$ was chosen arbitrarily, any optimal mechanism (\mathbf{p}, \mathbf{q}) should satisfy $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$ throughout \mathcal{T}_{IV} . This observation completes the proof. \square

The allocation probabilities given in Proposition 4 are satisfied by the favored-agent mechanism that assigns i^* as the favored agent and \bar{t}_{i^*} as the threshold. Furthermore, both type (i) and type (ii) version of this favored-agent mechanism satisfy the optimality condition in Proposition 3 so that they are both optimal when $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{i^*\}$. In our next result, we show that any other mechanism can only *weakly* Pareto robustly dominate the type (ii) variant of this favored-agent mechanism when $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{i^*\}$ is a singleton.

Proposition 5. *Assume that \mathcal{P} is a Markov ambiguity set of the form (2.6) and $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{i^*\}$ is a singleton. Denote by $(\mathbf{p}^*, \mathbf{q}^*)$ the allocation probabilities of the type (ii) favored-agent mechanism with the favored agent i^* and threshold $\nu^* = \bar{t}_{i^*}$. If a mechanism $(\mathbf{p}, \mathbf{q}) \in \mathcal{X}$ weakly Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$, then it satisfies*

$$\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) = \sum_{i \in \mathcal{I}} (p_i^*(\mathbf{t})t_i - q_i^*(\mathbf{t})c_i) \quad \forall \mathbf{t} \in \mathcal{T}.$$

Proposition 5 states that if a mechanism (\mathbf{p}, \mathbf{q}) weakly Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$, then mechanisms (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ earn the same payoff in all scenarios $\mathbf{t} \in \mathcal{T}$. This implies that no other mechanism can (strongly) Pareto robustly dominate $(\mathbf{p}^*, \mathbf{q}^*)$.

We sketch the proof idea for the two agents case detailed before. Recall that for this special case, we assume that $\underline{\mu}_2 < \underline{\mu}_1$ so that $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{1\}$ is a singleton and that $\bar{t}_2 > c_2 + \bar{t}_1$. We will again use the partition $\mathcal{T}_I - \mathcal{T}_V$ given in (2.8) and illustrated in Figure

2.2. In the following we first show that any mechanism (\mathbf{p}, \mathbf{q}) that weakly Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$ should be optimal. Hence, by Proposition 4, mechanisms (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ generate the same payoff throughout $\mathcal{T} \setminus \mathcal{T}_V$. Then, we will prove that the two mechanisms earn the same payoff also in \mathcal{T}_V .

To this end, fix a mechanism $(\mathbf{p}, \mathbf{q}) \in \mathcal{X}$ and assume that mechanism (\mathbf{p}, \mathbf{q}) weakly Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$, *i.e.*,

$$\mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} (p_i(\tilde{\mathbf{t}}) \tilde{t}_i - q_i(\tilde{\mathbf{t}}) c_i) \right] \geq \mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} (p_i^*(\tilde{\mathbf{t}}) \tilde{t}_i - q_i^*(\tilde{\mathbf{t}}) c_i) \right] \quad \forall \mathbb{P} \in \mathcal{P}.$$

By Theorem 3, $(\mathbf{p}^*, \mathbf{q}^*)$ is optimal in (MDP). As the expected payoff of mechanism (\mathbf{p}, \mathbf{q}) is at least as high as that of mechanism $(\mathbf{p}^*, \mathbf{q}^*)$ for any $\mathbb{P} \in \mathcal{P}$, mechanism (\mathbf{p}, \mathbf{q}) is also optimal. As $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{1\}$ is a singleton and as (\mathbf{p}, \mathbf{q}) is optimal, we know by Proposition 4 that (\mathbf{p}, \mathbf{q}) allocates the good to agent 1 without inspection if $t_2 - c_2 < \bar{t}_1$, *i.e.*, if $\mathbf{t} \in \mathcal{T} \setminus \mathcal{T}_V$. Thus, the allocation probabilities of the mechanisms (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ coincide on $\mathcal{T} \setminus \mathcal{T}_V$.

Next, consider any $\mathbf{t} \in \mathcal{T}_V$. Incentive compatibility ensures that $0 = p_2(t_1, t_2) \geq p_2(\mathbf{t}) - q_2(\mathbf{t})$, where the equality holds because $(t_1, t_2) \in \mathcal{T} \setminus \mathcal{T}_V$ and because we know that (\mathbf{p}, \mathbf{q}) allocates the good to agent 1 without inspection in $\mathcal{T} \setminus \mathcal{T}_V$. We may thus conclude that $p_2(\mathbf{t}) = q_2(\mathbf{t})$. Then, in scenario \mathbf{t} , the principal's payoff under (\mathbf{p}, \mathbf{q}) satisfies

$$\begin{aligned} \sum_{i \in \mathcal{I}} (p_i(\mathbf{t}) t_i - q_i(\mathbf{t}) c_i) &\leq p_2(\mathbf{t})(t_2 - c_2) + p_1(\mathbf{t}) t_1 \\ &\leq t_2 - c_2 = \sum_{i \in \mathcal{I}} (p_i^*(\mathbf{t}) t_i - q_i^*(\mathbf{t}) c_i), \end{aligned} \tag{2.10}$$

where the second inequality follows from the definition of \mathcal{T}_V that implies that $t_2 - c_2 \geq t_1$. The payoff of mechanism (\mathbf{p}, \mathbf{q}) therefore cannot exceed the one of $(\mathbf{p}^*, \mathbf{q}^*)$ throughout \mathcal{T}_V .

We will finally show that mechanisms (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ earn the same payoff on \mathcal{T}_V , *i.e.*, inequalities (2.10) hold as equalities. To this end, assume for the sake of contradiction that (\mathbf{p}, \mathbf{q}) earns a strictly lower payoff in scenario $\mathbf{t} \in \mathcal{T}_V$, *i.e.*,

$$\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) < t_2 - c_2 = \sum_{i \in \mathcal{I}} (p_i^*(\mathbf{t})t_i - q_i^*(\mathbf{t})c_i).$$

By Lemma 2, there exists $\hat{\mathbf{t}} \in \mathcal{T}$, where $\max_{i \neq i^*} \hat{t}_i < \hat{t}_{i^*}$, and $\mathbb{P} \in \mathcal{P}$ that satisfy: (i) $\mathbb{E}_{\mathbb{P}}[\tilde{t}_i] = \underline{\mu}_i \forall i \in \mathcal{I}$, (ii) $\mathbb{P}(\tilde{\mathbf{t}} \in \{\mathbf{t}, \hat{\mathbf{t}}\}) = 1$, (iii) $\mathbb{P}(\tilde{\mathbf{t}} = \mathbf{t}) > 0$. We already know from previous arguments that the payoff of mechanism (\mathbf{p}, \mathbf{q}) is lower than or equal to the payoff of $(\mathbf{p}^*, \mathbf{q}^*)$ in scenario $\hat{\mathbf{t}}$. In view of (ii) and (iii), the expected payoff earned by (\mathbf{p}, \mathbf{q}) is thus strictly lower than the one of $(\mathbf{p}^*, \mathbf{q}^*)$ under \mathbb{P} . As $\mathbb{P} \in \mathcal{P}$, this results in a contradiction with our initial assumption that (\mathbf{p}, \mathbf{q}) weakly Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$. Thus, mechanisms (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ earn the same payoff also on \mathcal{T}_V .

Proof of Proposition 5. We will again use the partition $\mathcal{T}_I - \mathcal{T}_V$ given in (2.9). Similarly to the sketch of the proof idea, we first show that (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ generate the same payoff throughout $\mathcal{T} \setminus \mathcal{T}_V$. Then, we will prove that the two mechanisms generate the same payoff also in \mathcal{T}_V .

To this end, fix a mechanism $(\mathbf{p}, \mathbf{q}) \in \mathcal{X}$ and assume that (\mathbf{p}, \mathbf{q}) weakly Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$. Mechanism (\mathbf{p}, \mathbf{q}) thus earns at least as high expected payoff as $(\mathbf{p}^*, \mathbf{q}^*)$ under every $\mathbb{P} \in \mathcal{P}$, *i.e.*, condition (2.1) holds. As $(\mathbf{p}^*, \mathbf{q}^*)$ is optimal by Theorem 3, this implies that (\mathbf{p}, \mathbf{q}) is also optimal in (MDP). As $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{i^*\}$ is a singleton, we thus know from Proposition 4 that (\mathbf{p}, \mathbf{q}) allocates the good to the favored agent i^* without inspection if $\max_{i \neq i^*} t_i - c_i < \bar{t}_{i^*}$, *i.e.*, if $\mathbf{t} \in \mathcal{T} \setminus \mathcal{T}_V$. Thus, the allocation probabilities of the mechanisms (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ coincide on $\mathcal{T} \setminus \mathcal{T}_V$, and they earn the same payoff throughout $\mathcal{T} \setminus \mathcal{T}_V$.

In the following we show that (\mathbf{p}, \mathbf{q}) can weakly Pareto robustly dominate $(\mathbf{p}^*, \mathbf{q}^*)$ only

if

$$\sum_{i \in \arg \max_{i' \in \mathcal{I}} t_{i'} - c_{i'}} p_i(\mathbf{t}) = 1 \quad \text{and} \quad p_i(\mathbf{t}) = q_i(\mathbf{t}) \quad \forall i \in \arg \max_{i' \in \mathcal{I}} t_{i'} - c_{i'} \quad (2.11)$$

for all $\mathbf{t} \in \mathcal{T}_V$. Note that (2.11) immediately implies that (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ generate the same payoff $\max_{i \in \mathcal{I}} t_i - c_i$ throughout \mathcal{T}_V .

Define now the set-valued function $\mathcal{I}(\mathbf{t}) = \{i \in \mathcal{I} \mid t_i \geq \bar{t}_{i^*}\}$ for $\mathbf{t} \in \mathcal{T}_V$. Note that $|\mathcal{I}(\mathbf{t})| \geq 1$ for all $\mathbf{t} \in \mathcal{T}_V$ thanks to the definition of \mathcal{T}_V , which ensures that there exists at least one agent $i \neq i^*$ with $t_i - c_i \geq \bar{t}_{i^*}$ and $\arg \max_{i \neq i^*} t_i - c_i \subseteq \mathcal{I}(\mathbf{t})$. We now prove by induction that (2.11) holds for all type profiles in $\mathcal{T}_V^n = \{\mathbf{t} \in \mathcal{T}_V \mid |\mathcal{I}(\mathbf{t})| = n\}$ for all $n \geq 1$.

As for the base step, set $n = 1$ and fix any $\mathbf{t} \in \mathcal{T}_V^1$. Thus, there exists exactly one agent $i^\circ \neq i^*$ such that $t_{i^\circ} \geq \bar{t}_{i^*}$. Incentive compatibility ensures that $p_{i^\circ}(\mathbf{t}) - q_{i^\circ}(\mathbf{t}) \leq p_{i^\circ}(\underline{t}_{i^\circ}, \mathbf{t}_{-i^\circ}) = 0$, where the equality holds because $(\underline{t}_{i^\circ}, \mathbf{t}_{-i^\circ}) \in \mathcal{T} \setminus \mathcal{T}_V$ and because (\mathbf{p}, \mathbf{q}) allocates the good to agent i^* without inspection on $\mathcal{T} \setminus \mathcal{T}_V$. We thus have $p_{i^\circ}(\mathbf{t}) = q_{i^\circ}(\mathbf{t})$. If $p_{i^\circ}(\mathbf{t}) < 1$, then

$$\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) \leq p_{i^\circ}(\mathbf{t})(t_{i^\circ} - c_{i^\circ}) + \sum_{i \neq i^\circ} p_i(\mathbf{t})t_i < \max_{i \in \mathcal{I}} t_i - c_i, \quad (2.12)$$

where the first inequality holds because $p_{i^\circ}(\mathbf{t}) = q_{i^\circ}(\mathbf{t})$ and $c_i > 0$ for all $i \neq i^\circ$. The second inequality follows from the assumption that $p_{i^\circ}(\mathbf{t}) < 1$ as well as the definition of \mathcal{T}_V^1 and the construction of i° , which imply that $t_{i^\circ} - c_{i^\circ} = \max_{i \in \mathcal{I}} t_i - c_i \geq \bar{t}_{i^*}$ and $\bar{t}_{i^*} > t_i$ for all $i \neq i^\circ$. This shows that (\mathbf{p}, \mathbf{q}) earns strictly less than $(\mathbf{p}^*, \mathbf{q}^*)$ in scenario \mathbf{t} . We next show that this fact contradicts inequality (2.1). Due to Lemma 2, there exists $\hat{\mathbf{t}} \in \mathcal{T}$, where $\max_{i \neq i^*} \hat{t}_i < \hat{t}_{i^*}$, and $\mathbb{P} \in \mathcal{P}$ that satisfy: (i) $\mathbb{E}_{\mathbb{P}}[\tilde{t}_i] = \underline{\mu}_i \forall i \in \mathcal{I}$, (ii) $\mathbb{P}(\tilde{\mathbf{t}} \in \{\mathbf{t}, \hat{\mathbf{t}}\}) = 1$,

(iii) $\mathbb{P}(\tilde{\mathbf{t}} = \mathbf{t}) > 0$. As $\hat{\mathbf{t}} \in \mathcal{T} \setminus \mathcal{T}_V$ by definition, we have

$$\begin{aligned} \mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} (p_i(\tilde{\mathbf{t}}) \tilde{t}_i - q_i(\tilde{\mathbf{t}}) c_i) \right] &= \alpha \sum_{i \in \mathcal{I}} (p_i(\mathbf{t}) t_i - q_i(\mathbf{t}) c_i) + (1 - \alpha) \sum_{i \in \mathcal{I}} (p_i(\hat{\mathbf{t}}) \hat{t}_i - q_i(\hat{\mathbf{t}}) c_i) \\ &< \alpha(t_{i^*} - c_{i^*}) + (1 - \alpha) \hat{t}_{i^*} = \mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} (p_i^*(\tilde{\mathbf{t}}) \tilde{t}_i - q_i^*(\tilde{\mathbf{t}}) c_i) \right], \end{aligned}$$

where $\alpha \in (0, 1]$ indicates the probability of $\tilde{\mathbf{t}} = \mathbf{t}$, and the inequality follows from (2.12) and the fact that the payoff at scenario $\hat{\mathbf{t}}$ is smaller than \hat{t}_{i^*} because $\max_{i \neq i^*} \hat{t}_i < \hat{t}_{i^*}$ and because (\mathbf{p}, \mathbf{q}) satisfies (FC) and $c_i > 0$ for all $i \in \mathcal{I}$. The strict inequality above implies that $(\mathbf{p}^*, \mathbf{q}^*)$ earns a strictly higher expected payoff than (\mathbf{p}, \mathbf{q}) under $\mathbb{P} \in \mathcal{P}$. It thus contradicts inequality (2.1) and our assumption that (\mathbf{p}, \mathbf{q}) weakly Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$. Hence, we have established (2.11) in scenario \mathbf{t} .

As for the induction step, assume that (2.11) holds throughout \mathcal{T}_V^n for some $n \geq 1$, and fix an arbitrary type profile $\mathbf{t} \in \mathcal{T}_V^{n+1}$. Thus, there exist exactly $n + 1$ agents i with types $t_i \geq \bar{t}_{i^*}$. For any agent $i \in \mathcal{I}(\mathbf{t}) \setminus \{i^*\}$ incentive compatibility dictates that $p_i(\mathbf{t}) - q_i(\mathbf{t}) \leq p_i(\underline{t}_i, \mathbf{t}_{-i}) = 0$, where the equality follows from Proposition 4 and the induction hypothesis because $(\underline{t}_i, \mathbf{t}_{-i}) \in \mathcal{T}_V^n \cup (\mathcal{T} \setminus \mathcal{T}_V)$. If $i^* \in \mathcal{I}(\mathbf{t})$, then we can make a similar argument for i^* . In fact, incentive compatibility dictates that $p_{i^*}(\mathbf{t}) - q_{i^*}(\mathbf{t}) \leq p_{i^*}(\underline{t}_{i^*}, \mathbf{t}_{-i^*}) = 0$, where the equality follows from the induction hypothesis because $(\underline{t}_{i^*}, \mathbf{t}_{-i^*}) \in \mathcal{T}_V^n$. In summary, we have thus shown that $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \mathcal{I}(\mathbf{t})$. The first condition in (2.11) can be proved by repeating the corresponding arguments from the base step almost verbatim. Details are omitted for brevity. We have thus established (2.11) in an arbitrary scenario $\mathbf{t} \in \mathcal{T}_V^{n+1}$. By induction, the revenues generated by the mechanisms (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ must therefore coincide throughout $\cup_{n \geq 1} \mathcal{T}_V^n = \mathcal{T}_V$. This observation completes the proof. \square

Proposition 5 shows that no other mechanism can Pareto robustly dominate the type (ii)

favored-agent mechanism with the favored agent i^* and threshold $\nu^* = \bar{t}_{i^*}$ and this mechanism is thus Pareto robustly optimal given that $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{i^*\}$ is a singleton. Next theorem proves that this Pareto robust optimality result continues to hold even when $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i$ is not a singleton.

Theorem 4. *If \mathcal{P} is equal to a Markov ambiguity set of the form (2.6), then any favored-agent mechanism of type (ii) with favored agent $i^* \in \arg \max_{i \in \mathcal{I}} \underline{\mu}_i$ and threshold value $\nu^* = \bar{t}_{i^*}$ is Pareto robustly optimal in (MDP).*

Proof. Let $(\mathbf{p}^*, \mathbf{q}^*)$ denote the allocation probabilities of the favored-agent mechanism described in Theorem 4. We know that $(\mathbf{p}^*, \mathbf{q}^*)$ is optimal from Theorem 3. To show that it is also Pareto robustly optimal, fix a mechanism $(\mathbf{p}, \mathbf{q}) \in \mathcal{X}$ and suppose that (\mathbf{p}, \mathbf{q}) weakly Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$, *i.e.*, condition (2.1) holds. We will show that (\mathbf{p}, \mathbf{q}) cannot (strictly) Pareto robustly dominate $(\mathbf{p}^*, \mathbf{q}^*)$.

If $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{i^*\}$ is a singleton, we know from Proposition 5 that (\mathbf{p}, \mathbf{q}) cannot generate strictly higher expected payoff under any $\mathbb{P} \in \mathcal{P}$, and $(\mathbf{p}^*, \mathbf{q}^*)$ is thus Pareto robustly optimal. Suppose now that $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i$ is not a singleton. Select any $\varepsilon \in (0, \bar{\mu}_{i^*} - \underline{\mu}_{i^*})$ that exists because $\underline{\mu}_{i^*} < \bar{\mu}_{i^*}$, and define

$$\mathcal{P}_\varepsilon = \{\mathbb{P} \in \mathcal{P} : \mathbb{E}_{\mathbb{P}}[\tilde{t}_{i^*}] \in [\underline{\mu}_{i^*} + \varepsilon, \bar{\mu}_{i^*}]\}.$$

Set \mathcal{P}_ε represents another Markov ambiguity set where the lowest mean value $\underline{\mu}_{i^*}$ of bidder i^* is shifted to $\underline{\mu}_{i^*} + \varepsilon$. Note that agent i^* becomes the unique agent with the maximum lowest mean value under \mathcal{P}_ε . As $\mathcal{P}_\varepsilon \subset \mathcal{P}$ by construction, we have

$$\mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} (p_i(\tilde{\mathbf{t}}) \tilde{t}_i - q_i(\tilde{\mathbf{t}}) c_i) \right] \geq \mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} (p_i^*(\tilde{\mathbf{t}}) \tilde{t}_i - q_i^*(\tilde{\mathbf{t}}) c_i) \right] \quad \forall \mathbb{P} \in \mathcal{P}_\varepsilon.$$

Thus, (\mathbf{p}, \mathbf{q}) also weakly Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$ under the Markov ambiguity set \mathcal{P}_ε . By Proposition 5, we can now conclude that (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ generate the same payoff for the principal in any scenario $\mathbf{t} \in \mathcal{T}$. This implies that the expected payoff of

(\mathbf{p}, \mathbf{q}) cannot exceed the one of $(\mathbf{p}^*, \mathbf{q}^*)$ under any distribution \mathbb{P} supported on \mathcal{T} . Thus, none of the inequalities in (2.1) can be strict, and (\mathbf{p}, \mathbf{q}) cannot Pareto robustly dominate $(\mathbf{p}^*, \mathbf{q}^*)$. The claim thus follows. \square

In the proof of Theorem 4, given a favored-agent mechanism $(\mathbf{p}^*, \mathbf{q}^*)$ of type (ii) with favored agent $i^* \in \arg \max_{i \in \mathcal{I}} \underline{\mu}_i$ and threshold value $\nu^* = \bar{t}_{i^*}$, we construct an auxiliary ambiguity set $\mathcal{P}_\epsilon \subseteq \mathcal{P}$ by increasing $\underline{\mu}_{i^*}$ to $\underline{\mu}_{i^*} + \epsilon$, where $\epsilon > 0$ is sufficiently small. By construction, agent i^* is the unique agent with the highest lower bound on the expected type under \mathcal{P}_ϵ . As $\mathcal{P}_\epsilon \subseteq \mathcal{P}$, any mechanism (\mathbf{p}, \mathbf{q}) that weakly Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$ under \mathcal{P} should also weakly Pareto robustly dominate $(\mathbf{p}^*, \mathbf{q}^*)$ under \mathcal{P}_ϵ . We then invoke Proposition 5 for \mathcal{P}_ϵ to conclude that (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ generate the same payoff in every scenario $\mathbf{t} \in \mathcal{T}$.

2.5 Markov Ambiguity Sets with Independent Types

Markov ambiguity sets studied in Section 2.4 contain distributions under which the agents' types are dependent. Throughout this section, we focus on a subset of the Markov ambiguity sets studied in Section 2.4 and assume that the agents' types are known to be mutually independent. In particular, we consider the Markov ambiguity set with independent types defined as

$$\mathcal{P} = \{\mathbb{P} \in \mathcal{P}_0(\mathcal{T}) \mid \mathbb{E}_{\mathbb{P}}[\tilde{t}_i] \in [\underline{\mu}_i, \bar{\mu}_i] \ \forall i \in \mathcal{I}, \ \tilde{t}_1, \dots, \tilde{t}_I \text{ are mutually independent under } \mathbb{P}\}. \quad (2.13)$$

As the Markov ambiguity set with independent types in (2.13) is a subset of the Markov ambiguity set in (2.6), the principal's optimal worst-case expected payoff cannot be lower than $\max_{i \in \mathcal{I}} \underline{\mu}_i$ in view of Proposition 2. The next proposition shows that the principal cannot improve her optimal worst-case expected payoff using the additional information of independence.

Proposition 6. *If \mathcal{P} is equal to a Markov ambiguity set of the form (2.13), then $z^* = \max_{i \in \mathcal{I}} \underline{\mu}_i$.*

The proof of Proposition 6 is identical to that of Proposition 2 as the agents' types are independent under the Dirac point mass at $\underline{\mu}$, which facilitates the proof.

Next theorem shows that there are again infinitely many optimal favored-agent mechanisms. The set of optimal favored-agent mechanisms characterized in Theorem 5 resembles to the one characterized in Theorem 3. Particularly, the selection criteria of a favored agent remains the same whereas the principal can select a lower threshold with the additional information of independence.

Theorem 5. *If \mathcal{P} is a Markov ambiguity set of the form (2.13), then any favored-agent mechanism with favored agent $i^* \in \arg \max_{i \in \mathcal{I}} \underline{\mu}_i$ and threshold value $\nu^* \geq \max_{i \in \mathcal{I}} \underline{\mu}_i$ is optimal in (MDP).*

Proof. Select any favored-agent mechanism with $i^* \in \arg \max_{i \in \mathcal{I}} \underline{\mu}_i$ and $\nu^* \geq \max_{i \in \mathcal{I}} \underline{\mu}_i$, denote by (\mathbf{p}, \mathbf{q}) its allocation probabilities. Recall first that this mechanism is feasible in (MDP). We will prove that (\mathbf{p}, \mathbf{q}) attains a worst-case expected payoff that is at least as large as $\max_{i \in \mathcal{I}} \underline{\mu}_i$, which implies via Proposition 6 that it is optimal in (MDP).

To this end, fix an arbitrary distribution $\mathbb{P} \in \mathcal{P}$ and suppose for ease of exposition that $\mathbb{P}(\max_{i \neq i^*} \tilde{t}_i - c_i < \nu^*)$, $\mathbb{P}(\max_{i \neq i^*} \tilde{t}_i - c_i = \nu^*)$ and $\mathbb{P}(\max_{i \neq i^*} \tilde{t}_i - c_i > \nu^*)$ are all strictly positive. We can write the principal's expected payoff from (\mathbf{p}, \mathbf{q}) under \mathbb{P} as

$$\begin{aligned} & \mathbb{P}\left(\max_{i \neq i^*} \tilde{t}_i - c_i < \nu^*\right) \mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} (p_i(\tilde{\mathbf{t}}) \tilde{t}_i - q_i(\tilde{\mathbf{t}}) c_i) \mid \max_{i \neq i^*} \tilde{t}_i - c_i < \nu^* \right] \\ & + \mathbb{P}\left(\max_{i \neq i^*} \tilde{t}_i - c_i = \nu^*\right) \mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} (p_i(\tilde{\mathbf{t}}) \tilde{t}_i - q_i(\tilde{\mathbf{t}}) c_i) \mid \max_{i \neq i^*} \tilde{t}_i - c_i = \nu^* \right] \\ & + \mathbb{P}\left(\max_{i \neq i^*} \tilde{t}_i - c_i > \nu^*\right) \mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} (p_i(\tilde{\mathbf{t}}) \tilde{t}_i - q_i(\tilde{\mathbf{t}}) c_i) \mid \max_{i \neq i^*} \tilde{t}_i - c_i > \nu^* \right]. \quad (2.14) \end{aligned}$$

If one or more of the terms $\mathbb{P}(\max_{i \neq i^*} \tilde{t}_i - c_i < \nu^*)$, $\mathbb{P}(\max_{i \neq i^*} \tilde{t}_i - c_i = \nu^*)$ and $\mathbb{P}(\max_{i \neq i^*} \tilde{t}_i - c_i > \nu^*)$ are zero, the equation (2.14) can be adjusted by removing the respective terms, and the proof proceeds similarly.

In the following we will show that all of the conditional expectations above and, therefore, the principal's expected payoff under \mathbb{P} are greater than or equal to $z^* = \max_{i \in \mathcal{I}} \underline{\mu}_i$. If $\max_{i \neq i^*} t_i - c_i < \nu^*$, condition (i) in Definition 3 implies that the principal's payoff amounts to t_{i^*} . This implies that

$$\begin{aligned} \mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} (p_i(\tilde{\mathbf{t}}) \tilde{t}_i - q_i(\tilde{\mathbf{t}}) c_i) \mid \max_{i \neq i^*} \tilde{t}_i - c_i < \nu^* \right] &= \mathbb{E}_{\mathbb{P}} \left[\tilde{t}_{i^*} \mid \max_{i \neq i^*} \tilde{t}_i - c_i < \nu^* \right] \\ &= \mathbb{E}_{\mathbb{P}} [\tilde{t}_{i^*}] = \mu_{i^*} \geq \max_{i \in \mathcal{I}} \underline{\mu}_i, \end{aligned}$$

where the second equality holds because the agents' types are independent. If $\max_{i \neq i^*} t_i - c_i > \nu^*$, then condition (ii) in Definition 3 implies that the principal's payoff amounts to $\max_{i \in \mathcal{I}} t_i - c_i$. We thus have

$$\begin{aligned} \mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} (p_i(\tilde{\mathbf{t}}) \tilde{t}_i - q_i(\tilde{\mathbf{t}}) c_i) \mid \max_{i \neq i^*} \tilde{t}_i - c_i > \nu^* \right] &= \mathbb{E}_{\mathbb{P}} \left[\max_{i \in \mathcal{I}} \tilde{t}_i - c_i \mid \max_{i \neq i^*} \tilde{t}_i - c_i > \nu^* \right] \\ &> \nu^* \geq \max_{i \in \mathcal{I}} \underline{\mu}_i. \end{aligned}$$

If $\max_{i \neq i^*} t_i - c_i = \nu^*$, then the allocation functions are defined either as in condition (i) or as in condition (ii) of Definition 3. If the allocation functions are defined as in condition (i), we have

$$\begin{aligned} \mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} (p_i(\tilde{\mathbf{t}}) \tilde{t}_i - q_i(\tilde{\mathbf{t}}) c_i) \mid \max_{i \neq i^*} \tilde{t}_i - c_i = \nu^* \right] &= \mathbb{E}_{\mathbb{P}} \left[\tilde{t}_{i^*} \mid \max_{i \neq i^*} \tilde{t}_i - c_i = \nu^* \right] \\ &= \mathbb{E}_{\mathbb{P}} [\tilde{t}_{i^*}] = \mu_{i^*} \geq \max_{i \in \mathcal{I}} \underline{\mu}_i, \end{aligned}$$

where the second equality again holds because the agents' types are independent. If the

allocation functions are defined as in condition (ii), on the other hand, then

$$\begin{aligned} \mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} (p_i(\tilde{\mathbf{t}}) \tilde{t}_i - q_i(\tilde{\mathbf{t}}) c_i) \mid \max_{i \neq i^*} \tilde{t}_i - c_i = \nu^* \right] &= \mathbb{E}_{\mathbb{P}} \left[\max_{i \in \mathcal{I}} \tilde{t}_i - c_i \mid \max_{i \neq i^*} \tilde{t}_i - c_i = \nu^* \right] \\ &\geq \nu^* \geq \max_{i \in \mathcal{I}} \underline{\mu}_i. \end{aligned}$$

In summary, we have shown that all of the conditional expectations in (2.14) have lower bound $\max_{i \in \mathcal{I}} \underline{\mu}_i$ and, therefore, the principal's expected payoff under \mathbb{P} is non-inferior to $z^* = \max_{i \in \mathcal{I}} \underline{\mu}_i$. As $\mathbb{P} \in \mathcal{P}$ was chosen arbitrarily, the principal's worst-case expected payoff is also non-inferior to z^* under the chosen favored-agent mechanism. Hence, by virtue of Proposition 6, any favored-agent mechanism with $i^* \in \arg \max_{i \in \mathcal{I}} \underline{\mu}_i$ and $\nu^* \geq \max_{i \in \mathcal{I}} \underline{\mu}_i$ is optimal in (MDP). \square

Similarly to previous sections, we next seek Pareto robustly optimal mechanisms for problem (MDP) under Markov ambiguity sets with independent types. To this end, we first present a few preliminary results some of which require the assumption that the set $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i$ is a singleton, *i.e.*, there is a single candidate for the favored agent. However, the Pareto robust optimality result of this section will not depend on this assumption.

Lemma 3. *If \mathcal{P} is a Markov ambiguity set of the form (2.13) and $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{i^*\}$ is a singleton, then, for any type profile $\mathbf{t} \in \mathcal{T}$ and any $\mu_{i^*} \in [\underline{\mu}_{i^*}, \bar{\mu}_{i^*}]$, there exist a scenario $\hat{\mathbf{t}} \in \mathcal{T}$, where $\max_{i \neq i^*} \hat{t}_i < \underline{\mu}_{i^*}$, and a discrete distribution $\mathbb{P} \in \mathcal{P}$ that satisfy the following properties: (i) $\mathbb{E}_{\mathbb{P}}[\tilde{t}_{i^*}] = \mu_{i^*}$, (ii) $\mathbb{P}(\tilde{t}_i \in \{t_i, \hat{t}_i\}) = 1$ for all $i \in \mathcal{I}$, (iii) $\mathbb{P}(\tilde{\mathbf{t}} = \mathbf{t}) > 0$.*

Proof. Consider arbitrary $\mathbf{t} \in \mathcal{T}$ and $\mu_{i^*} \in [\underline{\mu}_{i^*}, \bar{\mu}_{i^*}]$. We will construct a scenario $\hat{\mathbf{t}} \in \mathcal{T}$, where $\max_{i \neq i^*} \hat{t}_i < \underline{\mu}_{i^*}$, and a discrete distribution $\mathbb{P} \in \mathcal{P}$ that satisfies (i)–(iii). To this

end, we define \hat{t}_i through

$$\hat{t}_i = \begin{cases} t_i & \text{if } t_i = \underline{\mu}_i, \\ \underline{t}_i & \text{if } t_i > \underline{\mu}_i, \\ \underline{\mu}_i + \varepsilon & \text{if } t_i < \underline{\mu}_i, \end{cases} \quad \forall i \in \mathcal{I} \setminus \{i^*\} \quad \text{and} \quad \hat{t}_{i^*} = \begin{cases} t_{i^*} & \text{if } t_{i^*} = \mu_{i^*}, \\ \underline{t}_{i^*} & \text{if } t_{i^*} > \mu_{i^*}, \\ \mu_{i^*} + \varepsilon & \text{if } t_{i^*} < \mu_{i^*}, \end{cases}$$

where $\varepsilon \in (0, \min_{i \in \mathcal{I}} \bar{t}_i - \bar{\mu}_i) \cap (0, (\underline{\mu}_{i^*} - \max_{i \neq i^*} \underline{\mu}_i)/2)$ is a fixed positive number. Note that there exists such $\varepsilon > 0$ because $\underline{\mu}_i < \bar{\mu}_i < \bar{t}_i$ for all $i \in \mathcal{I}$ and $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{i^*\}$ is a singleton. We next show that $\hat{t}_i \in \mathcal{T}_i$ for all $i \in \mathcal{I}$ (i.e., $\hat{\mathbf{t}} \in \mathcal{T}$) and $\max_{i \neq i^*} \hat{t}_i < \underline{\mu}_{i^*}$. For any $i \in \mathcal{I}$, we have

$$\hat{t}_i \leq \bar{\mu}_i + \varepsilon \leq \bar{\mu}_i + \min_{j \in \mathcal{I}} (\bar{t}_j - \bar{\mu}_j) \leq \bar{\mu}_i + \bar{t}_i - \bar{\mu}_i = \bar{t}_i,$$

where the first inequality follows from the definition of \hat{t}_i , and the second inequality holds because $\varepsilon < \min_{j \in \mathcal{I}} \bar{t}_j - \bar{\mu}_j$. The definition of \hat{t}_i implies that we also have $\hat{t}_i \geq \underline{t}_i$. We thus showed that $\hat{\mathbf{t}} \in \mathcal{T}$.

For all $i \neq i^*$, we moreover have

$$\hat{t}_i \leq \underline{\mu}_i + \varepsilon \leq \underline{\mu}_i + (\underline{\mu}_{i^*} - \max_{j \neq i^*} \underline{\mu}_j)/2 \leq \underline{\mu}_i + (\underline{\mu}_{i^*} - \underline{\mu}_i)/2 < \underline{\mu}_{i^*},$$

where the first inequality again follows from the definition of \hat{t}_i , the second inequality holds because $\varepsilon < (\underline{\mu}_{i^*} - \max_{i \neq i^*} \underline{\mu}_i)/2$, and the fourth inequality holds because $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{i^*\}$ is a singleton. We thus showed that $\max_{i \neq i^*} \hat{t}_i < \underline{\mu}_{i^*}$.

Next, we will construct a discrete distribution \mathbb{P} through the marginal distributions $\mathbb{P}_i = \alpha_i \delta_{t_i} + (1 - \alpha_i) \delta_{\tilde{t}_i}$ of \tilde{t}_i 's, where $\alpha_i \in (0, 1]$ for all $i \in \mathcal{I}$. We will then show that \mathbb{P} belongs to the Markov ambiguity set \mathcal{P} and moreover satisfies the properties (i)–(iii). To this end,

we define α_i through

$$\alpha_i = \begin{cases} 1 & \text{if } t_i = \hat{t}_i, \\ (\underline{\mu}_i - \hat{t}_i)/(t_i - \hat{t}_i) & \text{if } t_i \neq \hat{t}_i, \end{cases} \quad \forall i \in \mathcal{I} \setminus \{i^*\},$$

and

$$\alpha_{i^*} = \begin{cases} 1 & \text{if } t_{i^*} = \hat{t}_{i^*}, \\ (\mu_{i^*} - \hat{t}_{i^*})/(t_{i^*} - \hat{t}_{i^*}) & \text{if } t_{i^*} \neq \hat{t}_{i^*}. \end{cases}$$

We first show that $\alpha_i \in (0, 1]$ for all $i \in \mathcal{I}$. For any $i \in \mathcal{I}$, it is sufficient to show that the claim holds if $t_i \neq \hat{t}_i$. For any $i \neq i^*$, if $t_i \neq \hat{t}_i$ and $t_i > \underline{\mu}_i$, we have

$$\alpha_i = (\underline{\mu}_i - \hat{t}_i)/(t_i - \hat{t}_i) = (\underline{\mu}_i - \underline{t}_i)/(t_i - \underline{t}_i) \in (0, 1),$$

where the second equality follows from the definition of \hat{t}_i , and the inclusion holds because $t_i > \underline{\mu}_i > \underline{t}_i$. If $t_i \neq \hat{t}_i$ and $t_i < \underline{\mu}_i$, on the other hand, we have $\alpha_i = -\varepsilon/(t_i - \underline{\mu}_i - \varepsilon) \in (0, 1)$, where the equality again follows from the definition of \hat{t}_i , and the inclusion holds because $t_i < \underline{\mu}_i < \underline{\mu}_i + \varepsilon$. Note that if $t_i = \underline{\mu}_i$, then $\hat{t}_i = t_i$ by definition, and $\alpha_i = 1$. One can similarly show that $\alpha_{i^*} \in (0, 1]$ by replacing $\underline{\mu}_{i^*}$ with μ_{i^*} in the above arguments. Thus, $\alpha_i \in (0, 1]$ for all $i \in \mathcal{I}$. We now define \mathbb{P} through the marginal distributions $\mathbb{P}_i = \alpha_i \delta_{t_i} + (1 - \alpha_i) \delta_{\hat{t}_i}$, $i \in \mathcal{I}$, as follows.

$$\mathbb{P}(\tilde{\mathbf{t}} = \mathbf{t}) = \prod_{i \in \mathcal{I}} \mathbb{P}_i(\tilde{t}_i = t_i) \quad \forall \mathbf{t} \in \mathcal{T}.$$

By construction, \tilde{t}_i 's are mutually independent under \mathbb{P} . Hence, the expected type of each $i \in \mathcal{I}$ amounts to $\mathbb{E}_{\mathbb{P}}[\tilde{t}_i] = \alpha_i t_i + (1 - \alpha_i) \hat{t}_i$.

We next show that $\mathbb{E}_{\mathbb{P}}[\tilde{t}_i] \in [\underline{\mu}_i, \bar{\mu}_i]$ for all $i \in \mathcal{I}$, which implies that $\mathbb{P} \in \mathcal{P}$. For any $i \neq i^*$, if $t_i = \hat{t}_i$, then we have $t_i = \hat{t}_i = \underline{\mu}_i$ by definition of \hat{t}_i . The expected type therefore

amounts to $\underline{\mu}_i$. If $t_i \neq \hat{t}_i$, on the other hand, we have

$$\mathbb{E}_{\mathbb{P}}[\tilde{t}_i] = \alpha_i t_i + (1 - \alpha_i) \hat{t}_i = \alpha_i (t_i - \hat{t}_i) + \hat{t}_i = \frac{\mu_i - \hat{t}_i}{t_i - \hat{t}_i} (t_i - \hat{t}_i) + \hat{t}_i = \underline{\mu}_i,$$

where the third equality follows from the definition of α_i . One can verify that $\mathbb{E}_{\mathbb{P}}[\tilde{t}_{i^*}] = \mu_{i^*}$ using similar arguments. We thus showed that $\mathbb{E}_{\mathbb{P}}[\tilde{t}_i] \in [\underline{\mu}_i, \bar{\mu}_i]$ for all $i \in \mathcal{I}$ and, therefore, $\mathbb{P} \in \mathcal{P}$.

It remains to show that \mathbb{P} satisfies (i)–(iii). As we have $\mathbb{E}_{\mathbb{P}}[\tilde{t}_{i^*}] = \mu_{i^*}$, property (i) holds. The definition of \mathbb{P} implies that (ii) and (iii) also hold. \square

The next technical lemma establishes a payoff equivalence result and will be used in the proof of the main Pareto robust optimality result of this section.

Lemma 4. *Assume that \mathcal{P} is equal to a Markov ambiguity set of the form (2.13), and $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{i^*\}$ is a singleton. Consider any subset $\mathcal{T}' = \prod_{i \in \mathcal{I}} \mathcal{T}'_i$ of \mathcal{T} such that (i) $\mathcal{T}'_i \supseteq \{t_i \in \mathcal{T}_i \mid t_i < \underline{\mu}_{i^*}\}$ for all $i \in \mathcal{I} \setminus \{i^*\}$ and (ii) either $\mathcal{T}'_{i^*} \subseteq [\underline{\mu}_{i^*}, \bar{\mu}_{i^*}]$ or $\mathcal{T}'_{i^*} = \mathcal{T}_{i^*}$. For any $(\mathbf{p}, \mathbf{q}), (\mathbf{p}', \mathbf{q}') \in \mathcal{X}$, if (\mathbf{p}, \mathbf{q}) weakly Pareto robustly dominates $(\mathbf{p}', \mathbf{q}')$ and*

$$\sum_{i \in \mathcal{I}} (p_i(\mathbf{t}) t_i - q_i(\mathbf{t}) c_i) \leq \sum_{i \in \mathcal{I}} (p'_i(\mathbf{t}) t_i - q'_i(\mathbf{t}) c_i) \quad \forall \mathbf{t} \in \mathcal{T}', \quad (2.15)$$

then (2.15) holds with equality.

Proof. Consider any subset $\mathcal{T}' = \prod_{i \in \mathcal{I}} \mathcal{T}'_i$ of \mathcal{T} such that (i) and (ii) holds. Also, consider any $(\mathbf{p}, \mathbf{q}), (\mathbf{p}', \mathbf{q}') \in \mathcal{X}$ such that (\mathbf{p}, \mathbf{q}) weakly Pareto robustly dominates $(\mathbf{p}', \mathbf{q}')$ and (2.15) holds. Suppose for the sake of contradiction that (2.15) is strict for some $\mathbf{t} \in \mathcal{T}'$.

We will characterize a discrete distribution $\mathbb{P} \in \mathcal{P}$ under which the expected payoff of mechanism (\mathbf{p}, \mathbf{q}) is strictly lower than that of $(\mathbf{p}', \mathbf{q}')$, which contradicts that (\mathbf{p}, \mathbf{q}) weakly Pareto robustly dominates $(\mathbf{p}', \mathbf{q}')$. By Lemma 3, for scenario \mathbf{t} and for any $\mu_{i^*} \in [\underline{\mu}_{i^*}, \bar{\mu}_{i^*}]$, there exist a scenario $\hat{\mathbf{t}} \in \mathcal{T}$, where $\max_{i \neq i^*} \hat{t}_i < \underline{\mu}_{i^*}$, and a discrete distribution

$\mathbb{P} \in \mathcal{P}$ that satisfy the following properties: (i) $\mathbb{E}_{\mathbb{P}}[\tilde{t}_{i^*}] = \mu_{i^*}$, (ii) $\mathbb{P}(\tilde{t}_i \in \{t_i, \hat{t}_i\}) = 1$ for all $i \in \mathcal{I}$, (iii) $\mathbb{P}(\tilde{\mathbf{t}} = \mathbf{t}) > 0$. We next show that there is always a $\mu_{i^*} \in [\underline{\mu}_{i^*}, \bar{\mu}_{i^*}]$ such that distribution \mathbb{P} also satisfies $\mathbb{P}(\tilde{\mathbf{t}} \in \mathcal{T}') = 1$. Note that if \mathbb{P} satisfies $\mathbb{P}(\tilde{t}_i \in \mathcal{T}'_i) = 1$ for all $i \in \mathcal{I}$, then it also satisfies $\mathbb{P}(\tilde{\mathbf{t}} \in \mathcal{T}') = 1$ as $\mathcal{T}' = \prod_{i \in \mathcal{I}} \mathcal{T}'_i$. First, suppose that $\mathcal{T}'_{i^*} \subseteq [\underline{\mu}_{i^*}, \bar{\mu}_{i^*}]$. For $\mu_{i^*} = t_{i^*} \in [\underline{\mu}_{i^*}, \bar{\mu}_{i^*}]$, properties (i)–(iii) on \mathbb{P} imply that if $\hat{t}_{i^*} \neq t_{i^*}$, then $\mathbb{P}(\tilde{t}_{i^*} = \hat{t}_{i^*}) = 0$. We thus have $\mathbb{P}(\tilde{t}_{i^*} = t_{i^*}) = 1$, which implies that $\mathbb{P}(\tilde{t}_{i^*} \in \mathcal{T}'_{i^*}) = 1$ as $t_{i^*} \in \mathcal{T}'_{i^*}$. For any $i \in \mathcal{I} \setminus \{i^*\}$, as $\hat{t}_i < \underline{\mu}_{i^*}$, we have $\hat{t}_i \in \mathcal{T}'_i \supseteq \{t_i \in \mathcal{T}_i \mid t_i < \underline{\mu}_{i^*}\}$ irrespective of the value of μ_{i^*} . Condition (ii) on \mathbb{P} thus implies that $\mathbb{P}(\tilde{t}_i \in \mathcal{T}'_i) = 1$. Suppose now that $\mathcal{T}'_{i^*} = \mathcal{T}_{i^*}$. Condition (ii) on \mathbb{P} implies that $\mathbb{P}(\tilde{t}_{i^*} \in \mathcal{T}'_{i^*}) = 1$ as $t_{i^*}, \hat{t}_{i^*} \in \mathcal{T}_{i^*} = \mathcal{T}'_{i^*}$. We already showed that $\mathbb{P}(\tilde{t}_i \in \mathcal{T}'_i) = 1$ for every other $i \in \mathcal{I} \setminus \{i^*\}$ irrespective of the value of μ_{i^*} . We can thus conclude that there always exists a $\mu_{i^*} \in [\underline{\mu}_{i^*}, \bar{\mu}_{i^*}]$ such that distribution \mathbb{P} from Lemma 3 also satisfies $\mathbb{P}(\tilde{\mathbf{t}} \in \mathcal{T}') = 1$.

Now, keeping in mind that \mathbb{P} is a discrete distribution with properties $\mathbb{P}(\tilde{\mathbf{t}} \in \mathcal{T}') = 1$ and $\mathbb{P}(\tilde{\mathbf{t}} = \mathbf{t}) > 0$, we can bound the principal's expected payoff from (\mathbf{p}, \mathbf{q}) under \mathbb{P} as follows:

$$\mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} (p_i(\tilde{\mathbf{t}}) \tilde{t}_i - q_i(\tilde{\mathbf{t}}) c_i) \right] < \mathbb{E}_{\mathbb{P}} \left[\sum_{i \in \mathcal{I}} (p'_i(\tilde{\mathbf{t}}) \tilde{t}_i - q'_i(\tilde{\mathbf{t}}) c_i) \right],$$

where the strict inequality follows from (2.15) and the assumption that (2.15) is strict for $\mathbf{t} \in \mathcal{T}'$. Therefore, we conclude that (\mathbf{p}, \mathbf{q}) cannot weakly Pareto robustly dominate $(\mathbf{p}', \mathbf{q}')$ unless the inequalities in (2.15) hold with equality. \square

Next proposition shows that the payoff equivalence result of Lemma 4 can be extended to the entire set \mathcal{T} of type profiles for a specific favored-agent mechanism.

Proposition 7. *Assume that \mathcal{P} is equal to a Markov ambiguity set of the form (2.13), and $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{i^*\}$ is a singleton. Denote by $(\mathbf{p}^*, \mathbf{q}^*)$ the allocation probabilities of the type (i) favored-agent mechanism with the favored agent i^* and threshold $\nu^* = \underline{\mu}_{i^*}$. If a*

mechanism $(\mathbf{p}, \mathbf{q}) \in \mathcal{X}$ weakly Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$, then it satisfies

$$\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) = \sum_{i \in \mathcal{I}} (p_i^*(\mathbf{t})t_i - q_i^*(\mathbf{t})c_i) \quad \forall \mathbf{t} \in \mathcal{T}.$$

We first sketch the proof idea focusing on a special case with two agents, where we assume that $\underline{\mu}_2 < \underline{\mu}_1$, *i.e.*, $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{1\}$, and that $\bar{t}_2 > c_2 + \underline{\mu}_1$. To this end, consider the following partition of the type space \mathcal{T} .

$$\begin{aligned} \mathcal{T}_I &= \{\mathbf{t} \in \mathcal{T} \mid t_1 \in (\underline{\mu}_1, \bar{\mu}_1] \text{ and } t_2 \leq \underline{\mu}_1\} \\ \mathcal{T}_{II} &= \{\mathbf{t} \in \mathcal{T} \mid t_1 \in (\underline{\mu}_1, \bar{\mu}_1], t_2 > \underline{\mu}_1 \text{ and } t_2 - c_2 \leq \underline{\mu}_1\} \\ \mathcal{T}_{II'} &= \{\mathbf{t} \in \mathcal{T} \mid t_1 \in (\underline{\mu}_1, \bar{\mu}_1], t_2 > \underline{\mu}_1 \text{ and } t_2 - c_2 > \underline{\mu}_1\} \\ \mathcal{T}_{III} &= \{\mathbf{t} \in \mathcal{T} \mid t_1 \notin (\underline{\mu}_1, \bar{\mu}_1] \text{ and } t_2 - c_2 \leq \underline{\mu}_1\} \\ \mathcal{T}_{III'} &= \{\mathbf{t} \in \mathcal{T} \mid t_1 \notin (\underline{\mu}_1, \bar{\mu}_1] \text{ and } t_2 - c_2 > \underline{\mu}_1\} \end{aligned}$$

Note that the condition $t_2 > \underline{\mu}_1$ in $\mathcal{T}_{II'}$ is redundant and given for ease of readability. We next show that we can replace the sets $\mathcal{T}_{II'}$ and $\mathcal{T}_{III'}$ with

$$\begin{aligned} \mathcal{T}_{IV} &= \{\mathbf{t} \in \mathcal{T} \mid t_1 = \underline{\mu}_1 \text{ and } t_2 - c_2 > \underline{\mu}_1\} \\ \mathcal{T}_V &= \{\mathbf{t} \in \mathcal{T} \mid t_1 \neq \underline{\mu}_1 \text{ and } t_2 - c_2 > \underline{\mu}_1\} \end{aligned}$$

and obtain a different partition of \mathcal{T} . To this end, first note that the intersection of \mathcal{T}_{IV} and \mathcal{T}_V is empty. Moreover, their union is given by $\{\mathbf{t} \in \mathcal{T} \mid t_2 - c_2 > \underline{\mu}_1\}$ that is the same as the union of $\mathcal{T}_{II'}$ and $\mathcal{T}_{III'}$. Thus, $\mathcal{T}_I, \mathcal{T}_{II}, \mathcal{T}_{III}, \mathcal{T}_{IV}, \mathcal{T}_V$ is a partition of the type space \mathcal{T} . We focus on this partition to simplify the arguments below. Sets $\mathcal{T}_I, \mathcal{T}_{II}, \mathcal{T}_{III}, \mathcal{T}_{IV}$ and \mathcal{T}_V are illustrated in Figure 2.3. Thanks to our standing assumptions, one can verify that all of these sets are nonempty. We emphasize that all simplifying assumptions will be relaxed in the formal proof of Proposition 7.

In the following, we show that any mechanism (\mathbf{p}, \mathbf{q}) that weakly Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$ generates the same payoff as $(\mathbf{p}^*, \mathbf{q}^*)$ in all scenarios $\mathbf{t} \in \mathcal{T}$. We will

prove this claim separately for each partition set and make use of Lemma 4. To this end, fix a mechanism $(\mathbf{p}, \mathbf{q}) \in \mathcal{X}$ and assume that mechanism (\mathbf{p}, \mathbf{q}) weakly Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$.

We first consider \mathcal{T}_I and note that it can be written as $\mathcal{T}_I = \mathcal{T}_{I1} \times \mathcal{T}_{I2} = (\underline{\mu}_1, \bar{\mu}_1] \times [t_2, \underline{\mu}_1]$. The principal's payoff under (\mathbf{p}, \mathbf{q}) in any $\mathbf{t} \in \mathcal{T}_I$ is given by $\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) \leq \sum_{i \in \mathcal{I}} p_i(\mathbf{t})t_i \leq t_1$, where the first inequality holds because $q_i(\mathbf{t})$ and c_i are non-negative, and the second inequality follows from (FC) and that $t_1 > t_2$ by definition of \mathcal{T}_I . As $(\mathbf{p}^*, \mathbf{q}^*)$ generates a payoff of t_1 in any $\mathbf{t} \in \mathcal{T}_I$ by definition, the payoff of $(\mathbf{p}^*, \mathbf{q}^*)$ is larger than or equal to the payoff of (\mathbf{p}, \mathbf{q}) in every $\mathbf{t} \in \mathcal{T}_I$. We assumed that (\mathbf{p}, \mathbf{q}) weakly Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$, and we showed that the payoff of $(\mathbf{p}^*, \mathbf{q}^*)$ cannot be lower than that of (\mathbf{p}, \mathbf{q}) in every $\mathbf{t} \in \mathcal{T}_I$. As $\mathcal{T}_I = \mathcal{T}_{I1} \times \mathcal{T}_{I2} = (\underline{\mu}_1, \bar{\mu}_1] \times [t_2, \underline{\mu}_1]$ satisfies the assumptions (i) and (ii) in Lemma 4, we can thus conclude that the payoffs of (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ coincide throughout \mathcal{T}_I by Lemma 4. Moreover, note that, for any $\mathbf{t} \in \mathcal{T}_I$, we have $t_2 < t_1$, $q_i(\mathbf{t}) \geq 0$, $c_i > 0$ and $\sum_{i \in \mathcal{I}} p_i(\mathbf{t}) \leq 1$. This implies that the payoff $\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i)$ of (\mathbf{p}, \mathbf{q}) can be t_1 only if $p_1(\mathbf{t}) = 1$ and $q_1(\mathbf{t}) = 0$.

Consider now any $\mathbf{t} \in \mathcal{T}_{II}$. Incentive compatibility ensures that $p_2(\mathbf{t}) - q_2(\mathbf{t}) \leq p_2(t_1, \underline{t}_2) = 0$, where the equality holds because $(t_1, \underline{t}_2) \in \mathcal{T}_I$ and because we know from before that (\mathbf{p}, \mathbf{q}) allocates the good to agent 1 without inspection throughout \mathcal{T}_I . As the allocation probabilities are non-negative and satisfy the (FC) condition $p_2(\mathbf{t}) \geq q_2(\mathbf{t})$, we thus have $p_2(\mathbf{t}) = q_2(\mathbf{t})$. This implies that the principal's payoff in \mathbf{t} satisfies $\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) \leq p_1(\mathbf{t})t_1 + p_2(\mathbf{t})(t_2 - c_2) \leq t_1$, where the first inequality holds




Figure 2.3: Partition of the type space \mathcal{T} , (Markov with Independence).

because $q_1(\mathbf{t})$ and c_1 are non-negative and $p_2(\mathbf{t}) = q_2(\mathbf{t})$, and the second inequality follows from (FC) and $t_2 - c_2 \leq \underline{\mu}_1 < t_1$. As $(\mathbf{p}^*, \mathbf{q}^*)$ generates a payoff of t_1 in any $\mathbf{t} \in \mathcal{T}_{II}$ by definition, the payoff of $(\mathbf{p}^*, \mathbf{q}^*)$ cannot be lower than that of (\mathbf{p}, \mathbf{q}) throughout \mathcal{T}_{II} . We next show that we can use Lemma 4 to conclude that the payoffs of two mechanisms should coincide throughout \mathcal{T}_{II} under our initial assumption, that is, (\mathbf{p}, \mathbf{q}) weakly Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$. To this end, note that $\mathcal{T}' = \mathcal{T}_I \cup \mathcal{T}_{II}$ can be written as $\mathcal{T}' = \mathcal{T}'_1 \times \mathcal{T}'_2 = (\underline{\mu}_1, \bar{\mu}_1] \times [t_2, c_2 + \underline{\mu}_1]$, and \mathcal{T}' satisfies the assumptions (i) and (ii) in Lemma 4. We showed that the payoff of $(\mathbf{p}^*, \mathbf{q}^*)$ is at least as high as the payoff of (\mathbf{p}, \mathbf{q}) in $\mathcal{T}' = \mathcal{T}_I \cup \mathcal{T}_{II}$. Thus, by Lemma 4, the payoff generated by (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ must coincide throughout $\mathbf{t} \in \mathcal{T}' = \mathcal{T}_I \cup \mathcal{T}_{II}$. Moreover, note that, for any $\mathbf{t} \in \mathcal{T}_{II}$, we have $t_2 - c_2 \leq \underline{\mu}_1 < t_1$, $q_i(\mathbf{t}) \geq 0$, $c_i > 0$ and $\sum_{i \in \mathcal{I}} p_i(\mathbf{t}) \leq 1$. This implies that the payoff $\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i)$ of (\mathbf{p}, \mathbf{q}) can be t_1 only if $p_1(\mathbf{t}) = 1$ and $q_1(\mathbf{t}) = 0$.

Next, consider any $\mathbf{t} \in \mathcal{T}_{III}$. Incentive compatibility ensures that $p_1(\mathbf{t}) \geq p_1(\bar{\mu}_1, t_2) - q_1(\bar{\mu}_1, t_2) = 1$, where the equality holds because $(\bar{\mu}_1, t_2) \in \mathcal{T}_I \cup \mathcal{T}_{II}$ and because we know from before that (\mathbf{p}, \mathbf{q}) allocates the good to agent 1 without inspection throughout $\mathcal{T}_I \cup \mathcal{T}_{II}$. Thus, the principal's payoff satisfies $\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) \leq p_1(\mathbf{t})t_1 - q_1(\mathbf{t})c_1 \leq t_1$, where the first inequality follows from $p_1(\mathbf{t}) = 1$, (FC) and non-negativity of \mathbf{q} , and the second inequality holds because $q_1(\mathbf{t})$ and c_1 are non-negative. As $(\mathbf{p}^*, \mathbf{q}^*)$ earns a payoff of t_1 in any $\mathbf{t} \in \mathcal{T}_{III}$ by definition, the payoff of $(\mathbf{p}^*, \mathbf{q}^*)$ thus cannot be lower than that of (\mathbf{p}, \mathbf{q}) throughout \mathcal{T}_{III} . Similarly to before, we next use Lemma 4 to show that the payoffs of the two mechanism should coincide throughout \mathcal{T}_{III} . To this end, note that $\mathcal{T}'' = \mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III}$ can be written as $\mathcal{T}'' = \mathcal{T}''_1 \times \mathcal{T}''_2 = \mathcal{T}_1 \times [t_2, c_2 + \underline{\mu}_1]$, and \mathcal{T}'' satisfies the assumptions (i) and (ii) in Lemma 4. We showed that the payoff of $(\mathbf{p}^*, \mathbf{q}^*)$ is at least as high as the payoff of (\mathbf{p}, \mathbf{q}) in $\mathcal{T}'' = \mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III}$. By Lemma 4, the payoffs of the two mechanisms must therefore coincide throughout $\mathcal{T}'' = \mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III}$. Moreover, for any $\mathbf{t} \in \mathcal{T}_{III}$, as $p_1(\mathbf{t}) = 1$ by incentive compatibility, $p_1(\mathbf{t}) \geq q_1(\mathbf{t}) \geq 0$, $c_1 > 0$ and $\sum_{i \in \mathcal{I}} p_i(\mathbf{t}) \leq 1$, mechanism (\mathbf{p}, \mathbf{q}) can generate a payoff of t_1 only if $p_1(\mathbf{t}) = 1$ and $q_1(\mathbf{t}) = 0$.

It remains to show that $(\mathbf{p}^*, \mathbf{q}^*)$ and (\mathbf{p}, \mathbf{q}) generate the same payoff in \mathcal{T}_{IV} and \mathcal{T}_V . To this end, we first show that agent 2, if allocated the good, should be inspected in any $\mathbf{t} \in \mathcal{T}_{IV} \cup \mathcal{T}_V$. For any such \mathbf{t} , the incentive compatibility ensures that $p_2(\mathbf{t}) - q_2(\mathbf{t}) \leq p_2(t_1, t_2) = 0$, where the equality holds because $(t_1, t_2) \in \mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III}$ and because (\mathbf{p}, \mathbf{q}) allocates the good to agent 1 without inspection throughout $\mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III}$. As the allocation probabilities are non-negative and satisfy the (FC) condition $p_2(\mathbf{t}) \geq q_2(\mathbf{t})$, we have $p_2(\mathbf{t}) = q_2(\mathbf{t})$.

Consider now any $\mathbf{t} \in \mathcal{T}_{IV}$. As $p_2(\mathbf{t}) = q_2(\mathbf{t})$, we have $\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) \leq p_1(\mathbf{t})t_1 + p_2(\mathbf{t})(t_2 - c_2) \leq t_2 - c_2$, where the first inequality holds because $q_1(\mathbf{t})$ and c_1 are non-negative and $p_2(\mathbf{t}) = q_2(\mathbf{t})$, and the second inequality follows from (FC) and $t_1 = \underline{\mu}_1 < t_2 - c_2$. As the payoff of $(\mathbf{p}^*, \mathbf{q}^*)$ is given by $t_2 - c_2$ in any $\mathbf{t} \in \mathcal{T}_{IV}$ by definition, the payoff of (\mathbf{p}, \mathbf{q}) cannot exceed that of $(\mathbf{p}^*, \mathbf{q}^*)$ throughout \mathcal{T}_{IV} . Consider the set $\mathcal{T}''' = \{\mathbf{t} \in \mathcal{T}_{III} \mid t_1 = \underline{\mu}_1\} \cup \mathcal{T}_{IV}$, which can be expressed as $\mathcal{T}''' = \{\underline{\mu}_1\} \times \mathcal{T}_2$, and note that $\mathcal{T}''' \subseteq \mathcal{T}_{III} \cup \mathcal{T}_{IV}$. Set \mathcal{T}''' satisfies the assumptions (i) and (ii) in Lemma 4, and the payoff of $(\mathbf{p}^*, \mathbf{q}^*)$ is at least as high as the one of (\mathbf{p}, \mathbf{q}) in \mathcal{T}''' . This implies, by Lemma 4, that the payoffs of the two mechanisms coincide throughout \mathcal{T}''' and therefore \mathcal{T}_{IV} . For any $\mathbf{t} \in \mathcal{T}_{IV}$, as $t_1 = \underline{\mu}_1 < t_2 - c_2$ and $p_2(\mathbf{t}) = q_2(\mathbf{t})$, (\mathbf{p}, \mathbf{q}) can generate a payoff of $t_2 - c_2$ only if $p_2(\mathbf{t}) = q_2(\mathbf{t}) = 1$.

Finally, consider any scenario $\mathbf{t} \in \mathcal{T}_V$. Incentive compatibility ensures that $p_1(\mathbf{t}) - q_1(\mathbf{t}) \leq p_1(\underline{\mu}_1, t_2) = 0$, where the equality holds because $(\underline{\mu}_1, t_2) \in \mathcal{T}_{IV}$ and because (\mathbf{p}, \mathbf{q}) allocates the good to agent 2 and inspects his report in \mathcal{T}_{IV} . As the allocation probabilities are non-negative and satisfy the (FC) condition $p_1(\mathbf{t}) \geq q_1(\mathbf{t})$, we may conclude that $p_1(\mathbf{t}) = q_1(\mathbf{t})$. Since we also have $p_2(\mathbf{t}) = q_2(\mathbf{t})$, we obtain $\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) \leq \sum_{i \in \mathcal{I}} p_i(\mathbf{t})(t_i - c_i) \leq \max_{i \in \mathcal{I}} t_i - c_i$, where the last inequality follows from the (FC) constraint $\sum_{i \in \mathcal{I}} p_i(\mathbf{t}) \leq 1$. As $(\mathbf{p}^*, \mathbf{q}^*)$ generates a payoff of $\max_{i \in \mathcal{I}} t_i - c_i$ in any $\mathbf{t} \in \mathcal{T}_V$ by definition, the payoff of (\mathbf{p}, \mathbf{q}) cannot exceed the payoff of $(\mathbf{p}^*, \mathbf{q}^*)$ throughout \mathcal{T}_V . Thus, (\mathbf{p}, \mathbf{q}) cannot generate a higher payoff than that of $(\mathbf{p}^*, \mathbf{q}^*)$ in $\mathcal{T} = \prod_{i \in \mathcal{I}} \mathcal{T}_i$, which trivially satisfies the assumptions (i) and (ii) in Lemma 4. This implies, by Lemma 4, that

the payoffs of the two mechanism must coincide throughout \mathcal{T} .

Proof of Proof of Proposition 7. Consider the following partition of the set \mathcal{T} .

$$\begin{aligned}\mathcal{T}_I &= \{\mathbf{t} \in \mathcal{T} \mid t_{i^*} \in (\underline{\mu}_{i^*}, \bar{\mu}_{i^*}] \text{ and } \max_{i \neq i^*} t_i \leq \underline{\mu}_{i^*}\} \\ \mathcal{T}_{II} &= \{\mathbf{t} \in \mathcal{T} \mid t_{i^*} \in (\underline{\mu}_{i^*}, \bar{\mu}_{i^*}] \text{ and } \max_{i \neq i^*} t_i > \underline{\mu}_{i^*} \text{ and } \max_{i \neq i^*} t_i - c_i \leq \underline{\mu}_{i^*}\} \\ \mathcal{T}_{II'} &= \{\mathbf{t} \in \mathcal{T} \mid t_{i^*} \in (\underline{\mu}_{i^*}, \bar{\mu}_{i^*}] \text{ and } \max_{i \neq i^*} t_i > \underline{\mu}_{i^*} \text{ and } \max_{i \neq i^*} t_i - c_i > \underline{\mu}_{i^*}\} \\ \mathcal{T}_{III} &= \{\mathbf{t} \in \mathcal{T} \mid t_{i^*} \notin (\underline{\mu}_{i^*}, \bar{\mu}_{i^*}] \text{ and } \max_{i \neq i^*} t_i - c_i \leq \underline{\mu}_{i^*}\} \\ \mathcal{T}_{III'} &= \{\mathbf{t} \in \mathcal{T} \mid t_{i^*} \notin (\underline{\mu}_{i^*}, \bar{\mu}_{i^*}] \text{ and } \max_{i \neq i^*} t_i - c_i > \underline{\mu}_{i^*}\}\end{aligned}$$

We can replace $\mathcal{T}_{II'}$ and $\mathcal{T}_{III'}$ with the following two sets to obtain a different partition of \mathcal{T} .

$$\begin{aligned}\mathcal{T}_{IV} &= \{\mathbf{t} \in \mathcal{T} \mid t_{i^*} = \underline{\mu}_{i^*} \text{ and } \max_{i \neq i^*} t_i - c_i > \underline{\mu}_{i^*}\} \\ \mathcal{T}_V &= \{\mathbf{t} \in \mathcal{T} \mid t_{i^*} \neq \underline{\mu}_{i^*} \text{ and } \max_{i \neq i^*} t_i - c_i > \underline{\mu}_{i^*}\}\end{aligned}$$

This is because \mathcal{T}_{IV} and \mathcal{T}_V are disjoint sets that have the same union as the union of $\mathcal{T}_{II'}$ and $\mathcal{T}_{III'}$. Throughout the proof we consider the partition $\mathcal{T}_I, \mathcal{T}_{II}, \mathcal{T}_{III}, \mathcal{T}_{IV}, \mathcal{T}_V$. Note that \mathcal{T}_I and \mathcal{T}_{III} are nonempty as $\arg \max_{i \in \mathcal{I}} \mu_i = \{i^*\}$ and $[\underline{\mu}_i, \bar{\mu}_i] \in (\underline{t}_i, \bar{t}_i)$ for all $i \in \mathcal{I}$, but sets $\mathcal{T}_{II}, \mathcal{T}_{IV}$ and \mathcal{T}_V can be empty if $\underline{\mu}_{i^*}$ or c_i for all $i \neq i^*$ are sufficiently large.

In the following, we inductively construct certain subsets of \mathcal{T} using the above partition. These subsets will have two important characteristics. First, they will be in the form of $\mathcal{T}' = \prod_{i \in \mathcal{I}} \mathcal{T}'_i$ and satisfy the assumptions (i) and (ii) in Lemma 4, *i.e.*, $\{t_i \in \mathcal{T}_i \mid t_i < \underline{\mu}_{i^*}\} \subseteq \mathcal{T}'_i$ for all $i \neq i^*$ and $\mathcal{T}'_{i^*} \subseteq [\underline{\mu}_{i^*}, \bar{\mu}_{i^*}]$ or $\mathcal{T}'_{i^*} = \mathcal{T}_{i^*}$. Second, for any mechanism $(\mathbf{p}, \mathbf{q}) \in \mathcal{X}$ that weakly Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$, we will inductively show that (\mathbf{p}, \mathbf{q}) generates a payoff lower than $(\mathbf{p}^*, \mathbf{q}^*)$ for all scenarios in the constructed subsets \mathcal{T}' . Thereafter, using Lemma 4, we will conclude that (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ generate the same payoff throughout \mathcal{T}' . The remainder of the proof is divided into five steps where the last step considers the set of all type profiles \mathcal{T} itself.

Step 1 (\mathcal{T}_I). For any $\mathbf{t} \in \mathcal{T}_I$, the principal's payoff under (\mathbf{p}, \mathbf{q}) satisfies

$$\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) \leq \sum_{i \in \mathcal{I}} p_i(\mathbf{t})t_i \leq t_{i^*} = \sum_{i \in \mathcal{I}} (p_i^*(\mathbf{t})t_i - q_i^*(\mathbf{t})c_i),$$

where the first inequality holds because $q_i(\mathbf{t})$ and c_i are non-negative, the second inequality follows from (FC) and that $\max_{i \neq i^*} t_i \leq \underline{\mu}_{i^*} < t_{i^*}$, and the equality follows from the definition of $(\mathbf{p}^*, \mathbf{q}^*)$. The payoff of $(\mathbf{p}^*, \mathbf{q}^*)$ is thus larger than or equal to the payoff of (\mathbf{p}, \mathbf{q}) in every $\mathbf{t} \in \mathcal{T}_I$. Moreover, note that \mathcal{T}_I can be written as $\prod_{i \in \mathcal{I}} \mathcal{T}_{I_i}$ where $\mathcal{T}_{I_{i^*}} = (\underline{\mu}_{i^*}, \bar{\mu}_{i^*}]$ and $\mathcal{T}_{I_i} = [\underline{t}_i, \underline{\mu}_{i^*}] \cap \mathcal{T}_i$ for all $i \neq i^*$. The set \mathcal{T}_I thus satisfies the assumptions (i) and (ii) in Lemma 4. By Lemma 4, we can thus conclude that the payoffs of (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ coincide in \mathcal{T}_I . In addition, note that, for any $\mathbf{t} \in \mathcal{T}_I$, we have $\max_{i \neq i^*} t_i < t_{i^*}$, $q_i(\mathbf{t}) \geq 0$, $c_i > 0$ and $\sum_{i \in \mathcal{I}} p_i(\mathbf{t}) \leq 1$. This implies that the payoff of (\mathbf{p}, \mathbf{q}) can match the payoff t_{i^*} of $(\mathbf{p}^*, \mathbf{q}^*)$ only if $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$.

Step 2 (\mathcal{T}_{II}). We will prove that if mechanism (\mathbf{p}, \mathbf{q}) weakly Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$ then it must satisfy $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$ for any $\mathbf{t} \in \mathcal{T}_{II}$, which implies that the payoff of (\mathbf{p}, \mathbf{q}) matches that of $(\mathbf{p}^*, \mathbf{q}^*)$ throughout \mathcal{T}_{II} . To this end, define the set-valued function $\mathcal{I}(\mathbf{t}) = \{i \in \mathcal{I} \mid t_i > \underline{\mu}_{i^*}\}$ for $\mathbf{t} \in \mathcal{T}_{II}$. Note that $|\mathcal{I}(\mathbf{t})| \geq 2$ for all $\mathbf{t} \in \mathcal{T}_{II}$ by the definition of \mathcal{T}_{II} , which ensures that $\max_{i \neq i^*} t_i > \underline{\mu}_{i^*}$ and $t_{i^*} \in (\underline{\mu}_{i^*}, \bar{\mu}_{i^*}]$. We now prove by induction that the claim holds in $\mathcal{T}_{II}^n = \{\mathbf{t} \in \mathcal{T}_{II} \mid |\mathcal{I}(\mathbf{t})| = n\}$ for all $n \geq 2$.

As for the base step, set $n = 2$ and fix any $\mathbf{t} \in \mathcal{T}_{II}^2$. Thus, there exists exactly one agent $i^\circ \neq i^*$ that satisfies $t_{i^\circ} > \underline{\mu}_{i^*}$. Incentive compatibility ensures that $p_{i^\circ}(\mathbf{t}) - q_{i^\circ}(\mathbf{t}) \leq p_{i^\circ}(\underline{t}_{i^\circ}, \mathbf{t}_{-i^\circ}) = 0$, where the equality holds because $(\underline{t}_{i^\circ}, \mathbf{t}_{-i^\circ}) \in \mathcal{T}_I$ and because we know from Step 1 that (\mathbf{p}, \mathbf{q}) allocates the good to agent i^* without inspection in \mathcal{T}_I . We thus have $p_{i^\circ}(\mathbf{t}) = q_{i^\circ}(\mathbf{t})$. Then, we have

$$\sum_{j \in \mathcal{I}} (p_j(\mathbf{t})t_j - q_j(\mathbf{t})c_j) \leq \sum_{j \neq i^\circ} p_j(\mathbf{t})t_j + p_{i^\circ}(\mathbf{t})(t_{i^\circ} - c_{i^\circ}) \leq t_{i^*} = \sum_{i \in \mathcal{I}} (p_i^*(\mathbf{t})t_i - q_i^*(\mathbf{t})c_i),$$

where the first inequality holds because $q_j(\mathbf{t})$ and c_j are non-negative and $p_{i^\circ}(\mathbf{t}) = q_{i^\circ}(\mathbf{t})$, the second inequality from (FC) and that $t_{i^\circ} - c_{i^\circ} \leq \underline{\mu}_{i^*} < t_{i^*}$ and $t_j \leq \underline{\mu}_{i^*}$ for all $j \in \mathcal{I} \setminus \{i^\circ, i^*\}$, and the equality follows from the definition of $(\mathbf{p}^*, \mathbf{q}^*)$. As scenario \mathbf{t} is chosen arbitrarily, the payoff of (\mathbf{p}, \mathbf{q}) thus cannot exceed that of $(\mathbf{p}^*, \mathbf{q}^*)$ throughout \mathcal{T}_{II}^2 . Recalling the conclusion from Step 1, we now know that this relation between the payoffs is true for the set $\mathcal{T}_I \cup \mathcal{T}_{II}^2$.

For any $i^\circ \in \mathcal{I} \setminus \{i^*\}$, define $\mathcal{T}_{II}^2(i^\circ)$ as the subset of \mathcal{T}_{II}^2 where i° is the only agent with type $t_{i^\circ} > \underline{\mu}_{i^*}$ and note that $\mathcal{T}_{II}^2 = \cup_{i^\circ \in \mathcal{I} \setminus \{i^*\}} \mathcal{T}_{II}^2(i^\circ)$. Consider an arbitrary $i^\circ \in \mathcal{I} \setminus \{i^*\}$ and the set $\mathcal{T}_I \cup \mathcal{T}_{II}^2(i^\circ)$, which can be written as $\mathcal{T}_I \cup \mathcal{T}_{II}^2(i^\circ) = \prod_{i \in \mathcal{I}} (\mathcal{T}_I \cup \mathcal{T}_{II}^2(i^\circ))_i$, where $(\mathcal{T}_I \cup \mathcal{T}_{II}^2(i^\circ))_{i^*} = (\underline{\mu}_{i^*}, \bar{\mu}_{i^*}]$, $(\mathcal{T}_I \cup \mathcal{T}_{II}^2(i^\circ))_{i^\circ} = [t_{i^\circ}, c_{i^\circ} + \underline{\mu}_{i^*}] \cap \mathcal{T}_{i^\circ}$ and $(\mathcal{T}_I \cup \mathcal{T}_{II}^2(i^\circ))_i = [t_i, \underline{\mu}_{i^*}] \cap \mathcal{T}_i$ for all $i \notin \{i^*, i^\circ\}$. The set $\mathcal{T}_I \cup \mathcal{T}_{II}^2(i^\circ)$ satisfies the assumptions (i) and (ii) in Lemma 4. Mechanisms (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ thus generate the same payoff throughout $\mathcal{T}_I \cup \mathcal{T}_{II}^2(i^\circ)$ by Lemma 4. By definition, the payoff of $(\mathbf{p}^*, \mathbf{q}^*)$ amounts to t_{i^*} in $\mathcal{T}_I \cup \mathcal{T}_{II}^2(i^\circ)$. For any $\mathbf{t} \in \mathcal{T}_{II}^2(i^\circ)$, as $t_{i^\circ} - c_{i^\circ} < t_{i^*}$, $t_i < t_{i^*}$ for all $i \notin \{i^*, i^\circ\}$ and $p_{i^\circ}(\mathbf{t}) = q_{i^\circ}(\mathbf{t})$, (\mathbf{p}, \mathbf{q}) can generate a payoff of t_{i^*} only if $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$. As i° is chosen arbitrarily, we have $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$ throughout \mathcal{T}_{II}^2 .

As for the induction step, assume that $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$ for all $\mathbf{t} \in \mathcal{T}_{II}^n$ and for some $n \geq 2$, and fix a scenario $\mathbf{t} \in \mathcal{T}_{II}^{n+1}$. Thus, there exists exactly $n + 1$ agents i that satisfy $t_i > \underline{\mu}_{i^*}$. For any agent $i \in \mathcal{I}(\mathbf{t}) \setminus \{i^*\}$, incentive compatibility dictates that $p_i(\mathbf{t}) - q_i(\mathbf{t}) \leq p_i(t_i, \mathbf{t}_{-i}) = 0$, where the equality follows from $(t_i, \mathbf{t}_{-i}) \in \mathcal{T}_{II}^n$ and the induction hypothesis. We thus have $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \mathcal{I}(\mathbf{t}) \setminus \{i^*\}$. Then,

$$\begin{aligned} \sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) &\leq \sum_{i \notin \mathcal{I}(\mathbf{t}) \setminus \{i^*\}} p_i(\mathbf{t})t_i + \sum_{i \in \mathcal{I}(\mathbf{t}) \setminus \{i^*\}} p_i(\mathbf{t})(t_i - c_i) \\ &\leq t_{i^*} = \sum_{i \in \mathcal{I}} (p_i^*(\mathbf{t})t_i - q_i^*(\mathbf{t})c_i), \end{aligned}$$

where the first inequality holds because $q_i(\mathbf{t})$ and c_i are non-negative and $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \mathcal{I}(\mathbf{t}) \setminus \{i^*\}$, the second inequality follows from (FC) and that $t_i - c_i \leq \underline{\mu}_{i^*}$ for all

$i \in \mathcal{I}(\mathbf{t}) \setminus \{i^*\}$ and $t_i \leq \underline{\mu}_{i^*} < t_{i^*}$ for $i \in \mathcal{I} \setminus \mathcal{I}(\mathbf{t})$, and the equality follows from the definition of $(\mathbf{p}^*, \mathbf{q}^*)$. As scenario \mathbf{t} is chosen arbitrarily, the payoff from (\mathbf{p}, \mathbf{q}) is thus less than or equal to that of $(\mathbf{p}^*, \mathbf{q}^*)$ throughout \mathcal{T}_{II}^{n+1} . By Step 1, this relationship between the payoffs holds true for the set $\mathcal{T}_I \cup \mathcal{T}_{II}^{n+1}$.

For any subset $\mathcal{I}' \ni i^*$ of agents with $|\mathcal{I}'| = n+1$, define $\mathcal{T}_{II}^{n+1}(\mathcal{I}')$ as the subset of \mathcal{T}_{II}^{n+1} where $t_i > \underline{\mu}_{i^*}$ for all $i \in \mathcal{I}'$. Note that the union of $\mathcal{T}_{II}^{n+1}(\mathcal{I}')$ over all $\mathcal{I}' \subseteq \mathcal{I}$ with $|\mathcal{I}'| = n+1$ and $i^* \in \mathcal{I}'$ gives us the set \mathcal{T}_{II}^{n+1} . Consider now an arbitrary $\mathcal{I}' \ni i^*$ with $|\mathcal{I}'| = n+1$ and the set $\mathcal{T}_I \cup \mathcal{T}_{II}^{n+1}(\mathcal{I}')$, which can be written as $\mathcal{T}_I \cup \mathcal{T}_{II}^{n+1}(\mathcal{I}') = \prod_{i \in \mathcal{I}} (\mathcal{T}_I \cup \mathcal{T}_{II}^{n+1}(\mathcal{I}'))_i$, where $(\mathcal{T}_I \cup \mathcal{T}_{II}^{n+1}(\mathcal{I}'))_{i^*} = (\underline{\mu}_{i^*}, \bar{\mu}_{i^*}]$, $(\mathcal{T}_I \cup \mathcal{T}_{II}^{n+1}(\mathcal{I}'))_i = [t_i, c_i + \underline{\mu}_{i^*}] \cap \mathcal{T}_i$ for $i \in \mathcal{I}' \setminus \{i^*\}$ and $(\mathcal{T}_I \cup \mathcal{T}_{II}^{n+1}(\mathcal{I}'))_i = [t_i, \underline{\mu}_{i^*}] \cap \mathcal{T}_i$ for all $i \in \mathcal{I} \setminus \mathcal{I}'$. The set $\mathcal{T}_I \cup \mathcal{T}_{II}^{n+1}(\mathcal{I}')$ satisfies the assumptions (i) and (ii) in Lemma 4. The payoffs of (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ thus coincide in $\mathcal{T}_I \cup \mathcal{T}_{II}^{n+1}(\mathcal{I}')$ by Lemma 4. By definition, the payoff of $(\mathbf{p}^*, \mathbf{q}^*)$ amounts to t_{i^*} throughout $\mathcal{T}_I \cup \mathcal{T}_{II}^{n+1}(\mathcal{I}')$. For any $\mathbf{t} \in \mathcal{T}_{II}^{n+1}(\mathcal{I}')$, as $t_i - c_i < t_{i^*}$ and $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \mathcal{I}' \setminus \{i^*\}$ and $t_i < t_{i^*}$ for all $i \in \mathcal{I} \setminus \mathcal{I}'$, mechanism (\mathbf{p}, \mathbf{q}) can generate a payoff of t_{i^*} only if $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$. As \mathcal{I}' is chosen arbitrarily, we have $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$ throughout \mathcal{T}_{II}^{n+1} . This thus completes the induction step.

In summary, the allocation probabilities of any mechanism (\mathbf{p}, \mathbf{q}) that weakly Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$ should satisfy $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$ throughout $\cup_{n \geq 2} \mathcal{T}_{II}^n = \mathcal{T}_{II}$.

Step 3 (\mathcal{T}_{III}). Next, fix any type profile $\mathbf{t} \in \mathcal{T}_{III}$. Incentive compatibility ensures that $p_{i^*}(\mathbf{t}) \geq p_{i^*}(\bar{\mu}_{i^*}, \mathbf{t}_{-i^*}) - q_{i^*}(\bar{\mu}_{i^*}, \mathbf{t}_{-i^*}) = 1$, where the equality holds because $(\bar{\mu}_{i^*}, \mathbf{t}_{-i^*}) \in \mathcal{T}_I \cup \mathcal{T}_{II}$ and because we know from Step 1 and 2 that (\mathbf{p}, \mathbf{q}) allocates the good to agent i^* without inspection in $\mathcal{T}_I \cup \mathcal{T}_{II}$. We thus have $p_{i^*}(\mathbf{t}) = 1$ and

$$\sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) \leq p_{i^*}(\mathbf{t})t_{i^*} + q_{i^*}(\mathbf{t})c_{i^*} \leq t_{i^*} = \sum_{i \in \mathcal{I}} (p_i^*(\mathbf{t})t_i - q_i^*(\mathbf{t})c_i),$$

where the first inequality follows from (FC) and non-negativity of $q_i(\mathbf{t})$, the second inequality holds because $q_{i^*}(\mathbf{t}) \geq 0$ and $c_{i^*} > 0$, and the equality follows from the definition of $(\mathbf{p}^*, \mathbf{q}^*)$. As scenario \mathbf{t} is chosen arbitrarily, (\mathbf{p}, \mathbf{q}) cannot generate a payoff higher than $(\mathbf{p}^*, \mathbf{q}^*)$ throughout \mathcal{T}_{III} . By Steps 1 and 2, this relation between the payoffs holds for the set $\mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III}$. Note that the set $\mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III}$ can be written as $\mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III} = \prod_{i \in \mathcal{I}} (\mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III})_i$, where $(\mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III})_{i^*} = \mathcal{T}_{i^*}$ and $(\mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III})_i = [\underline{t}_i, c_i + \underline{\mu}_{i^*}] \cap \mathcal{T}_i$ for all $i \in \mathcal{I} \setminus \{i^*\}$. The set $\mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III}$ thus satisfies the assumptions (i) and (ii) in Lemma 4. The payoffs of (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ thus coincide throughout $\mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III}$ by Lemma 4. As $q_{i^*}(\mathbf{t}) \geq 0$, $c_{i^*} > 0$ and (\mathbf{p}, \mathbf{q}) satisfies the (FC), mechanism (\mathbf{p}, \mathbf{q}) can generate a payoff of t_{i^*} in a scenario $\mathbf{t} \in \mathcal{T}_{III}$ only if $p_{i^*}(\mathbf{t}) = 1$ and $q_{i^*}(\mathbf{t}) = 0$.

Step 4 (\mathcal{T}_{IV}). In this step, we will show that any mechanism (\mathbf{p}, \mathbf{q}) that weakly Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$ must satisfy

$$\sum_{i \in \arg \max_{i' \in \mathcal{I}} t_{i'} - c_{i'}} p_i(\mathbf{t}) = 1 \quad \text{and} \quad p_i(\mathbf{t}) = q_i(\mathbf{t}) \quad \forall i \in \arg \max_{i' \in \mathcal{I}} t_{i'} - c_{i'} \quad (2.16)$$

for all $\mathbf{t} \in \mathcal{T}_{IV}$. This immediately implies that (\mathbf{p}, \mathbf{q}) generates the same payoff as $(\mathbf{p}^*, \mathbf{q}^*)$ throughout \mathcal{T}_{IV} . To this end, define the set-valued function $\mathcal{I}(\mathbf{t}) = \{i \in \mathcal{I} \mid t_i > \underline{\mu}_{i^*}\}$ for $\mathbf{t} \in \mathcal{T}_{IV}$. Note that $|\mathcal{I}(\mathbf{t})| \geq 1$ and $i^* \notin \mathcal{I}(\mathbf{t})$ for all $\mathbf{t} \in \mathcal{T}_{IV}$ thanks to the definition of \mathcal{T}_{IV} , which ensures that $\max_{i \neq i^*} t_i - c_i > \underline{\mu}_{i^*}$ and $t_{i^*} = \underline{\mu}_{i^*}$. We will prove by induction that (2.16) holds in $\mathcal{T}_{IV}^n = \{\mathbf{t} \in \mathcal{T}_{IV} \mid |\mathcal{I}(\mathbf{t})| = n\}$ for all $I - 1 \geq n \geq 1$.

As for the base step, set $n = 1$ and fix a scenario $\mathbf{t} \in \mathcal{T}_{IV}^1$. Thus, exactly one agent i° satisfies $t_{i^\circ} > \underline{\mu}_{i^*}$. Incentive compatibility ensures that $p_{i^\circ}(\mathbf{t}) - q_{i^\circ}(\mathbf{t}) \leq p_{i^\circ}(t_{i^\circ}, \mathbf{t}_{-i^\circ}) = 0$, where the equality follows from that $(t_{i^\circ}, \mathbf{t}_{-i^\circ}) \in \mathcal{T}_{III}$ and Step 3. We thus have $p_{i^\circ}(\mathbf{t}) =$

$q_{i^\circ}(\mathbf{t})$. Then,

$$\begin{aligned} \sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) &\leq \sum_{i \neq i^\circ} p_i(\mathbf{t})t_i + p_{i^\circ}(\mathbf{t})(t_{i^\circ} - c_{i^\circ}) \\ &\leq t_{i^\circ} - c_{i^\circ} = \sum_{i \in \mathcal{I}} (p_i^*(\mathbf{t})t_i - q_i^*(\mathbf{t})c_i), \end{aligned}$$

where the first inequality holds because $q_i(\mathbf{t})$ and c_i are non-negative and $p_{i^\circ}(\mathbf{t}) = q_{i^\circ}(\mathbf{t})$, the second inequality follows from (FC) and that $t_{i^\circ} - c_{i^\circ} > \underline{\mu}_{i^*}$ and $t_i \leq \underline{\mu}_{i^*}$ for all $i \in \mathcal{I} \setminus \{i^\circ\}$, and the equality follows from the definition of $(\mathbf{p}^*, \mathbf{q}^*)$. As scenario \mathbf{t} is chosen arbitrarily, (\mathbf{p}, \mathbf{q}) generates a payoff less than or equal to that of from $(\mathbf{p}^*, \mathbf{q}^*)$ throughout \mathcal{T}_{IV}^1 .

We now define the set $\mathcal{T}'_{III} = \{\mathbf{t} \in \mathcal{T}_{III} \mid t_{i^*} = \underline{\mu}_{i^*}\}$ that is a subset of \mathcal{T}_{III} . Recalling the findings in Step 3, the payoff of (\mathbf{p}, \mathbf{q}) cannot be higher than that of $(\mathbf{p}^*, \mathbf{q}^*)$ throughout $\mathcal{T}'_{III} \cup \mathcal{T}_{IV}^1$. For any $i^\circ \in \mathcal{I} \setminus \{i^*\}$, denote by $\mathcal{T}_{IV}^1(i^\circ)$ the subset of \mathcal{T}_{IV}^1 where i° is the only agent whose type $t_{i^\circ} > \underline{\mu}_{i^*}$ and note that $\cup_{i^\circ \in \mathcal{I} \setminus \{i^*\}} \mathcal{T}_{IV}^1(i^\circ) = \mathcal{T}_{IV}^1$. We have $\mathcal{T}'_{III} \cup \mathcal{T}_{IV}^1(i^\circ) = \prod_{i \in \mathcal{I}} (\mathcal{T}'_{III} \cup \mathcal{T}_{IV}^1(i^\circ))_i$, where $(\mathcal{T}'_{III} \cup \mathcal{T}_{IV}^1(i^\circ))_{i^*} = \{\underline{\mu}_{i^*}\}$, $(\mathcal{T}'_{III} \cup \mathcal{T}_{IV}^1(i^\circ))_{i^\circ} = \mathcal{T}_{i^\circ}$ and $(\mathcal{T}'_{III} \cup \mathcal{T}_{IV}^1(i^\circ))_i = [\underline{t}_i, \underline{\mu}_{i^*}] \cap \mathcal{T}_i$ for all $i \in \mathcal{I} \setminus \{i^*, i^\circ\}$. The set $\mathcal{T}'_{III} \cup \mathcal{T}_{IV}^1(i^\circ)$ satisfies the assumptions (i) and (ii) in Lemma 4, and the payoffs of (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ thus coincide throughout $\mathcal{T}'_{III} \cup \mathcal{T}_{IV}^1(i^\circ)$ by Lemma 4. As $t_{i^\circ} - c_{i^\circ} > t_{i^*} = \underline{\mu}_{i^*} \geq t_i$ for all $i \in \mathcal{I} \setminus \{i^\circ, i^*\}$ and $p_{i^\circ}(\mathbf{t}) = q_{i^\circ}(\mathbf{t})$, the payoff of (\mathbf{p}, \mathbf{q}) can match the payoff $\max_{i \in \mathcal{I}} t_i - c_i = t_{i^\circ} - c_{i^\circ}$ of $(\mathbf{p}^*, \mathbf{q}^*)$ only if $p_{i^\circ}(\mathbf{t}) = q_{i^\circ}(\mathbf{t}) = 1$. We thus established (2.16) in $\mathcal{T}_{IV}^1(i^\circ)$. As agent i° is chosen arbitrarily, the claim holds in \mathcal{T}_{IV}^1 .

As for the induction step, assume that (2.16) holds throughout \mathcal{T}_{IV}^n for some $n \geq 1$, and fix a scenario $\mathbf{t} \in \mathcal{T}_{IV}^{n+1}$. Thus, there exists exactly $n+1$ agents $i \neq i^*$ that satisfy $t_i > \underline{\mu}_{i^*}$. For any agent $i \in \mathcal{I}(\mathbf{t})$, incentive compatibility dictates that $p_i(\mathbf{t}) - q_i(\mathbf{t}) \leq p_i(\underline{t}_i, \mathbf{t}_{-i}) = 0$, where the equality follows from $(\underline{t}_i, \mathbf{t}_{-i}) \in \mathcal{T}_{III} \cup \mathcal{T}_{IV}^n$. Indeed, if $(\underline{t}_i, \mathbf{t}_{-i}) \in \mathcal{T}_{III}$, then agent $i^* \notin \mathcal{I}(\mathbf{t})$ receives the good so that $p_i(\underline{t}_i, \mathbf{t}_{-i}) = 0$, and if $(\underline{t}_i, \mathbf{t}_{-i}) \in \mathcal{T}_{IV}^n$, the equality follows from the induction hypothesis. We thus have $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \mathcal{I}(\mathbf{t})$

and, by the definition of \mathcal{T}_{IV} , in particular for all $i \in \arg \max_{j \in \mathcal{I}} t_j - c_j$. Then,

$$\begin{aligned} \sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) &\leq \sum_{i \notin \mathcal{I}(\mathbf{t})} p_i(\mathbf{t})t_i + \sum_{i \in \mathcal{I}(\mathbf{t})} p_i(\mathbf{t})(t_i - c_i) \\ &\leq \max_{i \in \mathcal{I}} t_i - c_i = \sum_{i \in \mathcal{I}} (p_i^*(\mathbf{t})t_i - q_i^*(\mathbf{t})c_i), \end{aligned}$$

where the first inequality holds because $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \mathcal{I}(\mathbf{t})$ and $q_i(\mathbf{t})$ and c_i are non-negative, the second inequality follows from (FC) and that $\max_{j \in \mathcal{I}} t_j - c_j > \underline{\mu}_{i^*} = t_{i^*} \geq t_i$ for all $i \notin \mathcal{I}(\mathbf{t})$, and the equality follows from the definition of $(\mathbf{p}^*, \mathbf{q}^*)$. Thus, the payoff of (\mathbf{p}, \mathbf{q}) is less than or equal to the payoff of $(\mathbf{p}^*, \mathbf{q}^*)$ in \mathcal{T}_{IV}^{n+1} .

For any subset $\mathcal{I}' \not\ni i^*$ of agents with $|\mathcal{I}'| = n + 1$, denote by $\mathcal{T}_{IV}^{n+1}(\mathcal{I}')$ the subset of \mathcal{T}_{IV}^{n+1} where $t_i > \underline{\mu}_{i^*}$ for all $i \in \mathcal{I}'$. Note that the union of $\mathcal{T}_{IV}^{n+1}(\mathcal{I}')$ over all $\mathcal{I}' \subset \mathcal{I}$ with $|\mathcal{I}'| = n + 1$ and $i^* \notin \mathcal{I}'$ gives us the set \mathcal{T}_{IV} . Consider an arbitrary such \mathcal{I}' and the set $\mathcal{T}'_{III} \cup \mathcal{T}_{IV}^{n+1}(\mathcal{I}')$, which can be written as $\mathcal{T}'_{III} \cup \mathcal{T}_{IV}^{n+1}(\mathcal{I}') = \prod_{i \in \mathcal{I}} (\mathcal{T}'_{III} \cup \mathcal{T}_{IV}^{n+1}(\mathcal{I}'))_i$, where $(\mathcal{T}'_{III} \cup \mathcal{T}_{IV}^{n+1}(\mathcal{I}'))_{i^*} = \{\underline{\mu}_{i^*}\}$, $(\mathcal{T}'_{III} \cup \mathcal{T}_{IV}^{n+1}(\mathcal{I}'))_i = \mathcal{T}_i$ for all $i \in \mathcal{I}'$ and $(\mathcal{T}'_{III} \cup \mathcal{T}_{IV}^{n+1}(\mathcal{I}'))_i = [t_i, \underline{\mu}_{i^*}] \cap \mathcal{T}_i$ for all $i \in \mathcal{I} \setminus \mathcal{I}' \cup \{i^*\}$. The set $\mathcal{T}'_{III} \cup \mathcal{T}_{IV}^{n+1}(\mathcal{I}')$ satisfies the assumptions (i) and (ii) in Lemma 4. By this lemma, (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ thus generate the same payoff throughout $\mathcal{T}'_{III} \cup \mathcal{T}_{IV}^{n+1}(\mathcal{I}')$. As $\max_{j \in \mathcal{I}} t_j - c_j > \underline{\mu}_{i^*} \geq t_i$ for all $i \notin \mathcal{I}'$ and $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \mathcal{I}'$, mechanism (\mathbf{p}, \mathbf{q}) can match the payoff $\max_{i \in \mathcal{I}} t_i - c_i$ of $(\mathbf{p}^*, \mathbf{q}^*)$ in $\mathcal{T}_{IV}^{n+1}(\mathcal{I}')$ only if (2.16) holds in $\mathcal{T}_{IV}^{n+1}(\mathcal{I}')$. As \mathcal{I}' is chosen arbitrarily, (2.16) holds throughout \mathcal{T}_{IV}^{n+1} . This observation completes the induction step.

Step 5 (\mathcal{T}_V). In analogy to Step 4, we will show that (2.16) holds for every $\mathbf{t} \in \mathcal{T}_V$. This immediately implies that (\mathbf{p}, \mathbf{q}) generates the same payoff as $(\mathbf{p}^*, \mathbf{q}^*)$ in \mathcal{T}_V and, consequently, throughout \mathcal{T} . To this end, define the set-valued function $\mathcal{I}(\mathbf{t}) = \{i \in \mathcal{I} \mid t_i > \underline{\mu}_{i^*}\}$ for $\mathbf{t} \in \mathcal{T}_V$. Note that $|\mathcal{I}(\mathbf{t})| \geq 1$ for any $\mathbf{t} \in \mathcal{T}_V$ thanks to the definition of \mathcal{T}_V , which implies that $\max_{i \neq i^*} t_i - c_i > \underline{\mu}_{i^*}$. We will prove by induction that (2.16) holds for all type profiles in $\mathcal{T}_V^n = \{\mathbf{t} \in \mathcal{T}_V \mid |\mathcal{I}(\mathbf{t}) \setminus \{i^*\}| = n\}$ for all $n \geq 1$. Note that in any $\mathbf{t} \in \mathcal{T}_V^n$ there are n agents, each of which is different from i^* , whose types exceed $\underline{\mu}_{i^*}$.

Agent i^* 's type may or may not take a value above $\underline{\mu}_{i^*}$.

As for the base step, set $n = 1$ and fix any scenario $\mathbf{t} \in \mathcal{T}_V^1$. Thus, there is exactly one agent $i^\circ \neq i^*$ that satisfy $t_{i^\circ} > \underline{\mu}_{i^*}$. Incentive compatibility ensures that $p_{i^\circ}(\mathbf{t}) - q_{i^\circ}(\mathbf{t}) \leq p_{i^\circ}(\underline{t}_{i^\circ}, \mathbf{t}_{-i^\circ}) = 0$, where the equality follows from that $(\underline{t}_{i^\circ}, \mathbf{t}_{-i^\circ}) \in \mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III}$ and from Steps 1, 2 and 3. Similarly for agent i^* , we have $p_{i^*}(\mathbf{t}) - q_{i^*}(\mathbf{t}) \leq p_{i^*}(\underline{\mu}_{i^*}, \mathbf{t}_{-i^*}) = 0$, where the equality follows from that $(\underline{\mu}_{i^*}, \mathbf{t}_{-i^*}) \in \mathcal{T}_{IV}$ and Step 4. We thus have $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \mathcal{I}(\mathbf{t})$ and i^* , which may or may not be an element of $\mathcal{I}(\mathbf{t})$. Then, we have

$$\begin{aligned} \sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) &\leq \sum_{i \in \mathcal{I} \setminus \mathcal{I}(\mathbf{t}) \cup \{i^*\}} p_i(\mathbf{t})t_i + \sum_{i \in \mathcal{I}(\mathbf{t}) \setminus \{i^*\}} p_i(\mathbf{t})(t_i - c_i) + p_{i^*}(\mathbf{t})(t_{i^*} - c_{i^*}) \\ &\leq \max_{i \in \mathcal{I}} t_i - c_i = \sum_{i \in \mathcal{I}} (p_i^*(\mathbf{t})t_i - q_i^*(\mathbf{t})c_i), \end{aligned}$$

where the first equality holds because $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \mathcal{I}(\mathbf{t}) \cup \{i^*\}$ and $q_i(\mathbf{t})$ and c_i are non-negative, the second inequality from (FC) and that $\max_{j \in \mathcal{I}} t_j - c_j = \max_{j \in \mathcal{I}(\mathbf{t}) \cup \{i^*\}} t_j - c_j > \underline{\mu}_{i^*} \geq t_i$ for all $i \in \mathcal{I} \setminus \mathcal{I}(\mathbf{t}) \cup \{i^*\}$, and the equality follows from the definition of $(\mathbf{p}^*, \mathbf{q}^*)$. Thus, (\mathbf{p}, \mathbf{q}) cannot generate a payoff higher than that of $(\mathbf{p}^*, \mathbf{q}^*)$ throughout \mathcal{T}_V^1 . Recalling the findings in Steps 1–4, this relation between the payoffs holds true for $\mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III} \cup \mathcal{T}_{IV} \cup \mathcal{T}_V^1$.

For any $i^\circ \in \mathcal{I} \setminus \{i^*\}$, denote by $\mathcal{T}_V^1(i^\circ)$ the subset of \mathcal{T}_V^1 where i° is the only agent among $\mathcal{I} \setminus \{i^*\}$ with type $t_{i^\circ} > c_{i^\circ} + \underline{\mu}_{i^*}$. Note that $\cup_{i^\circ \in \mathcal{I} \setminus \{i^*\}} \mathcal{T}_V^1(i^\circ) = \mathcal{T}_V^1$. Now for an arbitrary $i^\circ \in \mathcal{I} \setminus \{i^*\}$, recall the set $\mathcal{T}_{IV}^1(i^\circ)$ from Step 4 and consider the set $\mathcal{T}' = \mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III} \cup \mathcal{T}_{IV}^1(i^\circ) \cup \mathcal{T}_V^1(i^\circ)$, which can be written as $\mathcal{T}' = \prod_{i \in \mathcal{I}} (\mathcal{T}')_i$, where $(\mathcal{T}')_{i^*} = \mathcal{T}_{i^*}$, $(\mathcal{T}')_{i^\circ} = \mathcal{T}_{i^\circ}$ and $(\mathcal{T}')_i = [t_i, \underline{\mu}_{i^*}] \cap \mathcal{T}_i$ for all $i \in \mathcal{I} \setminus \mathcal{I}' \cup \{i^*\}$. As the payoff of (\mathbf{p}, \mathbf{q}) cannot be higher than that of $(\mathbf{p}^*, \mathbf{q}^*)$ throughout \mathcal{T}' and the set \mathcal{T}' satisfies the assumptions (i) and (ii) from Lemma 4, the payoffs from (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ must coincide throughout \mathcal{T}' . As $\max_{i \in \{i^\circ, i^*\}} t_i - c_i > \underline{\mu}_{i^*} \geq t_j$ for all $j \notin \{i^\circ, i^*\}$ and $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \{i^\circ, i^*\}$, mechanism (\mathbf{p}, \mathbf{q}) can match the payoff $\max_{i \in \{i^\circ, i^*\}} t_i - c_i$ from $(\mathbf{p}^*, \mathbf{q}^*)$ only if it satisfies the conditions in (2.16). As agent i° was chosen arbitrarily, we must

have revenue equivalence for \mathcal{T}_V^1 .

As for the induction step, assume that (2.16) holds throughout \mathcal{T}_V^n for some $n \geq 1$ and fix any scenario $\mathbf{t} \in \mathcal{T}_V^{n+1}$. For any agent $i \in \mathcal{I}(\mathbf{t})$, incentive compatibility implies that $p_i(\mathbf{t}) - q_i(\mathbf{t}) \leq p_i(\underline{t}_i, \mathbf{t}_{-i}) = 0$, where the equality holds due to $(\underline{t}_i, \mathbf{t}_{-i}) \in \mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III} \cup \mathcal{T}_V^n$. Indeed, if $(\underline{t}_i, \mathbf{t}_{-i}) \in \mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III}$, then we must have $i \neq i^*$ and $p_i(\underline{t}_i, \mathbf{t}_{-i}) = 0$ from Steps 1, 2 and 3, and if $(\underline{t}_i, \mathbf{t}_{-i}) \in \mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III} \cup \mathcal{T}_V^n$, then the equality follows from the induction hypothesis. We thus have $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \mathcal{I}(\mathbf{t})$, and by the definition of \mathcal{T}_V , in particular for all $i \in \arg \max_{j \in \mathcal{I}} t_j - c_j$. Then, the principal's payoff in \mathbf{t} can be written as:

$$\begin{aligned} \sum_{i \in \mathcal{I}} (p_i(\mathbf{t})t_i - q_i(\mathbf{t})c_i) &\leq \sum_{i \notin \mathcal{I}(\mathbf{t})} p_i(\mathbf{t})t_i + \sum_{i \in \mathcal{I}(\mathbf{t})} p_i(\mathbf{t})(t_i - c_i) \\ &\leq \max_{i \in \mathcal{I}} t_i - c_i = \sum_{i \in \mathcal{I}} (p_i^*(\mathbf{t})t_i - q_i^*(\mathbf{t})c_i), \end{aligned}$$

where the first inequality follows because $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \mathcal{I}(\mathbf{t})$ and $c_i > 0$ for all $i \notin \mathcal{I}(\mathbf{t})$. The second inequality holds because the two sums represent a weighted average of $t_i - c_i$ for $i \in \mathcal{I}(\mathbf{t})$ and t_i for $i \notin \mathcal{I}(\mathbf{t})$. All these terms are smaller or equal to $\max_{i \in \mathcal{I}} t_i - c_i$. In particular, the definition of \mathcal{T}_V ensures that $\max_{j \in \mathcal{I}} t_j - c_j > \underline{\mu}_{i^*} \geq t_i$ for all $i \notin \mathcal{I}(\mathbf{t})$. This reasoning shows that the payoff from (\mathbf{p}, \mathbf{q}) cannot be higher than $(\mathbf{p}^*, \mathbf{q}^*)$ in \mathcal{T}_V^{n+1} .

For any subset $\mathcal{I}' \ni i^*$ of agents with $|\mathcal{I}' \setminus \{i^*\}| = n + 1$, denote by $\mathcal{T}_V^{n+1}(\mathcal{I}')$ the subset of \mathcal{T}_V^{n+1} where $t_i > \underline{\mu}_{i^*}$ for all $i \in \mathcal{I}'$. Note that the union of $\mathcal{T}_V^{n+1}(\mathcal{I}')$ over all $\mathcal{I}' \subset \mathcal{I}$ with $|\mathcal{I}'| = n + 1$ and $i^* \in \mathcal{I}'$ gives us the set \mathcal{T}_V . For an arbitrary such \mathcal{I}' , recall the set $\mathcal{T}_{IV}^{n+1}(\mathcal{I}' \setminus \{i^*\})$ from Step 4 and consider the set $\mathcal{T}' = \mathcal{T}_I \cup \mathcal{T}_{II} \cup \mathcal{T}_{III} \cup \mathcal{T}_{IV}^{n+1}(\mathcal{I}' \setminus \{i^*\}) \cup \mathcal{T}_V^{n+1}(\mathcal{I}')$, which can be written as $\mathcal{T}' = \prod_{i \in \mathcal{I}} (\mathcal{T}')_i$, where $(\mathcal{T}')_i = \mathcal{T}_i$ for all $i \in \mathcal{I}'$ and $(\mathcal{T}')_i = [\underline{t}_i, \underline{\mu}_{i^*}] \cap \mathcal{T}_i$ for all $i \in \mathcal{I} \setminus \mathcal{I}'$. The set \mathcal{T}' satisfies the assumptions (i) and (ii) in Lemma 4 so that the payoffs from (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ must coincide under the subset of scenarios \mathcal{T}' . As $\max_{i \in \mathcal{I}} t_i - c_i > \underline{\mu}_{i^*} \geq t_j$ for all $j \notin \mathcal{I}'$ and $p_i(\mathbf{t}) = q_i(\mathbf{t})$ for all $i \in \mathcal{I}'$,

mechanism (\mathbf{p}, \mathbf{q}) can match the payoff $\max_{i \in \mathcal{I}} t_i - c_i$ only if (2.16) holds in $\mathcal{T}_V^{n+1}(\mathcal{I}')$. As \mathcal{I}' was chosen arbitrarily, (2.16) holds throughout \mathcal{T}_V^{n+1} . This observation completes the induction step.

The above reasoning show that the principal's payoff under (\mathbf{p}, \mathbf{q}) and $(\mathbf{p}^*, \mathbf{q}^*)$ coincide in fact throughout the entire type space \mathcal{T} . Therefore, the proof is complete. \square

When $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{i^*\}$, Pareto robust optimality of the type (i) favored-agent mechanism $(\mathbf{p}^*, \mathbf{q}^*)$ with the favored agent i^* and threshold $\nu^* = \underline{\mu}_{i^*}$ follows from Proposition 7. That is, any mechanism (\mathbf{p}, \mathbf{q}) that weakly Pareto robustly dominates $(\mathbf{p}^*, \mathbf{q}^*)$ generates the same expected payoff for the principal under every distribution $\mathbb{P} \in \mathcal{P}$. Hence, $(\mathbf{p}^*, \mathbf{q}^*)$ cannot be Pareto robustly dominated when $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i = \{i^*\}$. Moreover, as for the general Markov ambiguity sets, we can show that $(\mathbf{p}^*, \mathbf{q}^*)$ remains to be Pareto robust optimal even if $\arg \max_{i \in \mathcal{I}} \underline{\mu}_i$ is not a singleton. We can do this by revising the proof of Theorem 4 in a way that only employs the distributions \mathbb{P} under which the agents' types are independent.

Theorem 6. *Assume that \mathcal{P} is equal to a Markov ambiguity set of the form (2.6) and the agents' types are independent. Then, the favored-agent mechanism of type (i) with favored agent $i^* \in \arg \max_{i \in \mathcal{I}} \underline{\mu}_i$ and threshold value $\nu^* = \max_{i \in \mathcal{I}} \underline{\mu}_i$ is Pareto robustly optimal in (MDP).*

Finally, let us consider Markov ambiguity sets in general and contrast the results of dependent and independent types settings. Given a favored agent $i^* \in \arg \max_{i \in \mathcal{I}} \underline{\mu}_i$, the principal should pick a threshold no less than \bar{t}_{i^*} to achieve optimality when the agents' types can be dependent. On the other hand, when the agents' types are independent, optimality can be achieved by setting the threshold to $\nu^* = \max_{i \in \mathcal{I}} \underline{\mu}_i$. Hence, using the independent types information, the principal can decrease the threshold from \bar{t}_{i^*} to $\underline{\mu}_{i^*}$ while retaining a worst-case expected payoff of $\max_{i \in \mathcal{I}} \underline{\mu}_i$. That is, the principal no longer needs to choose a threshold that is dependent on the choice of favored agent.

Chapter 3

Nonexclusive Competition for a Freelancer under Adverse Selection

3.1 Introduction

Consider a freelancer who has limited working hours either due to legal obligations (*e.g.*, 48 hrs/week) or natural constraints (*e.g.*, 24 hrs/day) and can serve multiple parties by allocating his time accordingly.¹ Suppose the freelancer values the leisure time that he can spare from his working hours. Hence, working an extra minute gets more costly as the allocated time for work gets higher (convex cost). On the other side of the market, multiple parties can benefit from the services of the freelancer but have limited information regarding the quality of the service (adverse selection). Furthermore, no buyer can pose limits on the freelancer regarding the contract deals made with the other buyers (nonexclusivity). In modern labor markets, nonexclusivity becomes more and more the rule. A

¹The EU's Working Time Directive (2003/88/EC) requires EU Member States to enforce a limit to weekly working hours: the average working time for each seven day period must not exceed 48 hours.

real-life example is a consultant who faces multiple firms seeking his expertise. What kind of trades shall we expect to arise in such a setup? ²

In this chapter, we characterize the equilibrium trades for this problem under the following setting: There are at least two buyers interested in the services of the freelancer. The freelancer has private information regarding the quality of his service that can be either low or high. The buyers share a common prior regarding the quality of the service provided by the freelancer. The buyers have linear preferences for quality and compete through offering contracts that specify a quantity (number of working hours) and a transfer (payment to the freelancer).³ The freelancer observes the offers and chooses the contracts that maximize his payoff. The preferences of (each type of) the freelancer are quasilinear: They are linear in the aggregate payment and display strictly increasing convex cost in the aggregate quantity.

In this context, we characterize the freelancer's aggregate trades in any pure-strategy equilibria. Our results can be summarized as follows: We provide two distinct conditions either of which is sufficient for the existence of a pure-strategy equilibrium. These conditions are also necessary so that there is no pure-strategy equilibrium if both fail to hold. Furthermore, they depend only on the preferences of the high-type freelancer:

- (i) At the no-trade point, the high-type freelancer is *not* willing to trade any amount of his time in exchange for a price equal to the average quality of the service.

²A recent study jointly conducted by Upwork (a global freelancing platform) and Freelancers Union highlights the growing share of the freelancers in the U.S. labor market. For instance, at the release date of the study, the contribution of freelance income to the U.S. economy is reported to be nearly \$1 trillion, which is almost 5% of U.S. GDP. Another important finding is that 45% of freelancers provide skilled services such as programming, marketing, IT, and business consulting. Hence, nearly half of the freelancers are offering their expertise in a steadily growing market that constitutes an important part of the U.S. economy. Source: <https://www.cnn.com/2019/10/03/skilled-freelancers-earn-more-per-hour-than-70percent-of-workers-in-us.html>. Retrieved on 2020-08-20.

³In this chapter, we focus explicitly on a labor market setting where multiple parties are interested in hiring a freelancer for a service that they cannot provide themselves. Because quantity traded is generally measured as time in labor markets, we restrict ourselves to the case where the freelancer trades only non-negative quantities.

- (ii) At any feasible trade point, the high-type freelancer is willing to trade any amount of his time in exchange for a price equal to the average quality of the service.

If condition (i) holds, then the high-type freelancer does not trade in equilibrium while the aggregate trade of the low-type depends on his preferences. In such equilibria, the buyers make zero profit from each of their traded contracts. On the other hand, if condition (ii) holds, then both types trade at the capacity and, there is cross-subsidization in equilibrium. In all of these equilibria, aggregate equilibrium trades are unique.

Our results contribute to the literature on competition under adverse selection. There are two classical papers in this literature: Akerlof [24] considers a market where the sellers are privately informed about the quality of their goods. The goods are non-divisible, and all trades take place at the same price. Because uninformed buyers do not consider trading at a price above the average quality of the goods, sellers of high-quality goods end up not trading in equilibrium. On the other hand, Rothschild & Stiglitz [25] considers a similar setup where uninformed buyers compete through contract offers for a divisible good. By offering different quantities at different unit prices, the buyers can screen the quality of the goods through sellers' contract choices. Rothschild & Stiglitz [25] allow only for exclusive competition, *i.e.*, each seller can only trade with at most one buyer. They show that, when an equilibrium exists, low-quality sellers trade efficiently while high-quality sellers trade a non-zero, but sub-optimal quantity.⁴

Attar et al. [1, 2] are the first to bring nonexclusive competition together with adverse selection. They observe that in many real-life market situations sellers simultaneously and secretly trade with several buyers. In their words, “nonexclusivity is the rule rather than the exception” in many markets. This is also true for the modern labor markets: many firms are simultaneously and secretly seeking the expertise of a freelancer. Hence, our

⁴Mas-Colell et al. [26, Chapter 13] provides an analysis of competitive labor markets under adverse selection as well. In Section B, they apply the model of [24] to the labor market, whereas in Section D, they consider the exclusive competition approach presented in [25]. Our setting differs from theirs in that the competition is nonexclusive.

work is complementary to Attar et al. [1] and Attar et al. [2]. These two papers differ from our work in two dimensions concerning the seller (the freelancer in our setup): (i) capacity constraint and (ii) convex cost. In [1], they consider a seller with a linear cost and a capacity constraint, whereas in [2], they consider a seller who has convex preferences but does not have any capacity constraints.⁵ Our model differs from [1] in that the freelancer has a convex cost in the aggregate quantity traded, and differs from [2] in that the freelancer is subject to a capacity constraint. Therefore, by bringing the capacity constraint and convex cost together, not only do we consider a natural and relevant setup for labor markets but also we provide a bridge between the results of [1] and [2]. Table 3.1 summarizes the differences between our work and these two papers.⁶

	This Chapter	Attar et al. [1]	Attar et al. [2]
Freelancer's Preferences	Quasilinear with strictly convex cost	Linear	Strictly quasiconcave
Capacity Constraint	✓	✓	✗
The quality of the service	High or Low	Continuum, discrete, or mixed	High or Low
Existence of equilibrium	Exists iff the high type is either not willing to trade at a price equal to the average quality at no-trade or willing to trade at a price equal to the average quality at any feasible trade	Always exists (for a large class of type distributions)	Exists iff the high type is not willing to trade at a price equal to the average quality at no-trade

⁵Even though [2] allow trades to be unrestricted in sign, they elaborate on the necessary changes on their results when only non-negative quantities can be sold.

⁶The comparison in Table 3.1 is provided only for the case of non-negative trades in [2].

Cross-subsidization	Depends on the preferences of the high-type	Possible in equilibrium	Ruled out in equilibrium
Aggregate Equilibrium Trades	Unique (both types trade at the capacity or the high type does not trade while the low type either trades efficiently, trades all of his capacity, or does not trade)	Unique (if the quality is low enough the freelancer trades all of his capacity or does not trade)	Unique (the high type does not trade, the low type either trades efficiently or does not trade at all)

Table 3.1: **Comparison with Attar et al. [1, 2]**

Our results confirm that the Akerlof-like equilibrium outcomes presented in the earlier works extend to our setting. For instance, if the freelancer with the high-quality service is not willing to work at a price equal to the expected quality, then an equilibrium can be supported where only the low-quality freelancer has a chance to trade. In this case, the buyers protect themselves against the information asymmetry by offering a contract that is only acceptable to the low-quality freelancer. On the other hand, when the high-quality freelancer is willing to work for the price equal to the expected quality, the buyers find it profitable to offer a pooling contract. Then, no equilibrium can be supported unless the capacity constraint is low enough in the sense that at any feasible trade point, the marginal cost of the high type is less than the expected quality. In this case, a pooling equilibrium exists, and both types trade at the capacity. In any equilibrium, competition pushes the price up so that the buyers end up having zero expected profit.

Inderst & Wambach [27, 28] study the role of capacity constraints in competitive screening models, which features an environment à la Rothschild & Stiglitz [25]. Somewhat parallel to our conclusion, they show that the presence of a capacity constraint alleviates the problem of non-existence of equilibrium in pure strategies. Their work departs from ours in two main aspects: They do not allow nonexclusive competition and assume

that each contract issuer (buyer in our setting) faces a capacity constraint. Such a constraint on the buyers' side limits their ability to unilaterally deviate and make a profit. In our setting, the capacity constraint is on the freelancer's side, which, together with the non-negative trades assumption, limits the set of feasible deviations available to the buyers.

3.2 The Model

A freelancer (seller) faces contract offers from multiple parties (buyers) who seek his services. He can serve more than one customer by allocating his working hours accordingly but, the number of working hours available to the freelancer is limited and denoted by \bar{Q}_C . He privately knows the quality of his service, which can be either H or L . The probability that the quality is of type i is commonly known to be $m_i \in (0, 1)$ for $i \in \{H, L\}$. That is, $m_H + m_L = 1$. The freelancer only cares about the aggregate hours he works, Q , and the aggregate monetary transfers he receives T . We assume that the freelancer has quasilinear preferences: The numerical representation of his payoff is $u_i(Q, T) = T - c_i(Q)$ where the cost function, c_i for $i \in \{H, L\}$, is a continuously differentiable, strictly convex real function defined over $[0, \bar{Q}_C]$. Hence, type i 's marginal rate of substitution of working hours for money is equal to his marginal cost, $c'_i(Q)$. We assume that for the same level of aggregate working hours, type H incurs a strictly higher marginal cost than type L . That is, $c'_H(Q) > c'_L(Q)$ for all $Q \in [0, \bar{Q}_C]$.

On the other side of the market, there are $n \geq 2$ identical buyers. Each buyer k offers a set of contracts, $C^k \subset \mathbb{R}^2$ consisting of required working hours and transfer bundles, that is denoted by (q, t) .⁷ We have $(0, 0) \in C^i$ for all buyers so that the freelancer may choose not to trade with any particular buyer. Each buyer only cares about his trade with

⁷As noted by Attar et al. [1, 2], we do not need to consider more general mechanisms in our setup. See [29] and [30] for further details.

the freelancer. Upon agreeing on a contract (q, t) with type i , a buyer earns a profit of $\nu_i q - t$ where ν_i is the (constant) marginal benefit of being served by type i . We assume that the marginal benefit from working with type H is strictly higher than type L . Hence, the expected quality of the service that is denoted by $\nu = m_H \nu_H + m_L \nu_L$ satisfies $\nu_H > \nu > \nu_L$.

Upon receiving the set of offers, the freelancer chooses a contract from each of the offered set of menus. Thus, type i freelancer needs to solve the following maximization problem:

$$\max \left\{ \sum_l t^l - c_i \left(\sum_l q^l \right) : \sum_l q^l \leq \bar{Q}_C, (q^l, t^l) \in C^l \text{ for each } l \right\}.$$

Menus of contracts are assumed to be compact so that, this problem always has a solution. We use the perfect Bayesian equilibrium concept and focus on pure-strategy equilibria where each type i of the freelancer chooses to trade a contract, (q_i^k, t_i^k) from the menus of contracts offered by each buyer k . Aggregate equilibrium trades for type i is denoted by $(Q_i, T_i) = (\sum_l q_i^l, \sum_l t_i^l)$. We define the indirect utility function that gives the maximum payoff that type i freelancer can achieve while trading a contract (q, t) with buyer k as follows:

$$\begin{aligned} z_i^{-k}(q, t) &= \max_{(q^l, t^l) \in C^l} t + \sum_{l \neq k} t^l - c_i \left(q + \sum_{l \neq k} q^l \right) \\ \text{s.t.} \quad &\sum_{l \neq k} q^l \leq \bar{Q}_C - q, \end{aligned}$$

In equilibrium, one should have $U_i = u_i(Q_i, T_i) = z_i^{-k}(q_i^k, t_i^k)$ for all i and k . As noted by Attar et al. [2], z_i^{-k} defined above may have discontinuities due to the capacity constraint. Therefore, the proofs in [2] that exploit the continuity of the indirect utility function are no longer valid in our setup. Yet, using the linearity of the freelancer's return on transfers, we can construct simple feasibility arguments to determine the conditions under which pure strategy equilibrium exists.

3.3 Equilibrium Characterization

After observing the contract offers, the freelancer solves the corresponding maximization problem considering the bilateral trades between each buyer and himself. Whenever an equilibrium exists, no buyer should be able to change his contract offer and increase his expected payoff. We derive properties of the equilibria which survive well-chosen buyer deviations following a similar methodology to that of Attar et al. [2], which parallels the solution methodology in Rothschild & Stiglitz [25]. Under nonexclusive competition, any buyer can build his deviation on the contracts offered by the other buyers. Considering type i freelancer's optimal choice, (Q_i, T_i) in aggregate, buyer k can fix arbitrary contracts from other buyer's menus which amounts to (Q^{-k}, T^{-k}) and deviate by offering $(q, t) = (Q_i - Q^{-k}, T_i - T^{-k})$. The first result in [2] derives equilibrium conditions on such a contract. In our problem setting, the same conditions hold for the feasible set of deviations even if the indirect utility function of the freelancer is discontinuous. In that vein, Lemma 5 shows that if some buyer k can improve his profits with type i , then his deviation should be traded by both types of the freelancer and, it should not be profitable in expectation. We define b_i^k as the profit of the buyer k from his trade with the type i freelancer. That is, $b_i^k = \nu_i q_i^k - t_i^k$. Similarly, we define expected profit of buyer k as $b^k = m_L b_L^k + m_H b_H^k$.

Lemma 5. *In equilibrium, for all $q \in [0, \bar{Q}_C]$ and t , if the freelancer can trade $(Q_i - q, T_i - t)$ with buyers other than k , then*

$$\nu_i q - t > b_i^k \quad \text{implies} \quad \nu q - t \leq b^k.$$

Proof. Assume that the freelancer can trade $(Q_H - q, T_H - t)$ with buyers other than k and $\nu_H q - t > b_H^k$ holds (the proof for type L is similar).

Consider the following deviation for buyer k : $\{(0, 0), (q, t + \epsilon_H), (q_L^k, t_L^k + \epsilon_L)\}$ for $\epsilon_H > \epsilon_L > 0$. Then, by trading $(q, t + \epsilon_H)$ with buyer k and trading $(Q_H - q, T_H - t)$ with the buyers other than k , type H can strictly increase his payoff after the deviation.

Therefore, type H strictly prefers trading $(q, t + \epsilon_H)$ to trading $(0, 0)$ with buyer k . Now, fix arbitrary contracts from the menus offered by buyers other than k and assume that they amount to (Q^{-k}, T^{-k}) . We know that in equilibrium $U_H \geq u_H(Q^{-k} + q_L^k, T^{-k} + t_L^k)$ holds for all $Q^{-k} + q_L^k$ that is less than or equal to the capacity. Since the payoff of the freelancer is linear in transfers, we have $u_H(Q_H, T_H + \epsilon_H) > U_H + \epsilon_L \geq u_H(Q^{-k} + q_L^k, T^{-k} + t_L^k + \epsilon_L)$ for all feasible $Q^{-k} + q_L^k$. Hence, type H freelancer also strictly prefers trading $(q, t + \epsilon_H)$ to $(q_L^k, t_L^k + \epsilon_L)$. On the other hand, type L can strictly increase his profits by trading $(q_L^k, t_L^k + \epsilon_L)$. Thus, type L strictly prefers trading $(q_L^k, t_L^k + \epsilon_L)$ with buyer k to trading $(0, 0)$ with buyer k . Assume that type L trades $(q_L^k, t_L^k + \epsilon_L)$ after the deviation. Then, buyer k earns:

$$m_H(\nu_H q - t) + m_L b_L^k - (m_H \epsilon_H + m_L \epsilon_L).$$

This is strictly greater than b^k for small enough ϵ_H and ϵ_L . Hence, in equilibrium, type L should also trade $(q, t + \epsilon_H)$, and the resulting profit for buyer k cannot be higher than b^k in equilibrium:

$$\nu q - t - m_H \epsilon_H \leq b^k.$$

The result follows by letting ϵ_H approach zero. □

In the proof of Lemma 5, we consider a deviation for buyer k in which he offers three contracts: $(0, 0)$, (q, t) that improves buyer k 's profits with type i , and the equilibrium contract traded with the other type. Note that monetary transfers of the last two contracts are increased by a small margin so that their respective types prefer them to no-trade contracts. If the indirect utility function is continuous as in Attar et al. [2], then one can choose the increment value for the last contract in a way that type i chooses to trade (q, t) after the deviation. Although the function $z_i^k(q, t)$ may be discontinuous in our setting, we make use of the quasi-linear nature of u_i to design the deviation contracts in the same manner. If the other type chooses to trade his equilibrium contract, then this is a profitable deviation for buyer k . Hence, after the deviation, both types should trade (q, t) , and the resulting payoff should be less than or equal to the equilibrium payoff of buyer k .

Now, consider the payoff of each type in terms of aggregate equilibrium trades. We can write the following two inequalities:

$$\begin{aligned} T_L - c_L(Q_L) &\geq T_H - c_L(Q_H), \\ T_H - c_H(Q_H) &\geq T_L - c_H(Q_L). \end{aligned}$$

Since the cost function is continuously differentiable, summing up the above inequalities and employing the fundamental theorem of calculus leads to the following:

$$\begin{aligned} c_H(Q_L) - c_H(Q_H) &\geq c_L(Q_L) - c_L(Q_H), \\ \int_{Q_H}^{Q_L} c'_H(x) dx &\geq \int_{Q_H}^{Q_L} c'_L(x) dx. \end{aligned}$$

Hence, due to the assumption $c'_H(Q) > c'_L(Q)$, type L should provide a higher level of service in any equilibrium, *i.e.*, $Q_L \geq Q_H$.

Let S_L be the aggregate profit of the buyers gained from additionally trading $(Q_L - Q_H, T_L - T_H)$ with type L freelancer, *i.e.*, $S_L = \nu_L(Q_L - Q_H) - (T_L - T_H)$; and similarly, $S_H = \nu_H(Q_H - Q_L) - (T_H - T_L)$. Attar et al. [2] prove that in equilibrium, $S_L \leq 0$, and the expected payoff of each buyer is zero. Using these results, they characterize the candidate equilibria. These conclusions remain to be true in our setting because their proofs rely only on Lemma 5. Before going further on, we define the aggregate profit of the buyers from type i freelancer as $B_i = \sum_l b_i^l$ for both $i \in \{H, L\}$. Hence, the aggregate expected profit of the buyers is $B = \sum_l b^l = \sum_l (m_L b_L^l + m_H b_H^l) = m_L B_L + m_H B_H$.

Proposition 8. (Attar et al. [2]) *In any equilibrium, $S_L \leq 0$ and $B = 0$ so that $b^l = 0$ for each l . Moreover, the following statements hold.*

- (i) *In any pooling equilibrium, $T_L = \nu Q_L = T_H = \nu Q_H$.*
- (ii) *In any separating equilibrium, $Q_L > Q_H \geq 0$ holds with $T_H = \nu Q_H$ and $T_L - T_H = \nu_L(Q_L - Q_H)$.*

As a consequence of Proposition 8, we obtain the following immediate result, which will be useful in the remaining proofs.

Corollary 1. *In any equilibrium, $B_H \geq 0 \geq B_L$ and $S_L = 0$ hold.*

In words, the aggregate profit of the buyers gained from additionally trading $(Q_L - Q_H, T_L - T_H)$ with the low-type freelancer is always zero in equilibrium. Furthermore, the profit from the high-type freelancer subsidizes for the loss from the low-type freelancer if an equilibrium exhibits cross-subsidization.

Next, we derive conditions on the set of equilibria in which buyers make a strictly positive aggregate profit with the high-type freelancer. Lemma 6 below shows that, in such equilibria, the marginal cost of the high-type freelancer should be less than or equal to the expected quality of the service. Furthermore, the contract offers of any single buyer should not be essential for the aggregate equilibrium trades of the high-type freelancer. In other words, the high-type freelancer should be able to trade at the same aggregate level even if a buyer withdraws his offers:

Lemma 6. *If in equilibrium $B_H > 0$, then*

$$c'_H(Q_H) \begin{cases} = \nu & \text{if } Q_H < \bar{Q}_C, \\ \leq \nu & \text{if } Q_H = \bar{Q}_C. \end{cases}$$

Moreover, for each buyer k , the freelancer can trade (Q_H, T_H) with buyers other than k .

Proof. The result for $c'_H(Q_H) = \nu$ when $B_H > 0$ is due to Lemma 3 of Attar et al. [2].⁸ Suppose to the contrary that $B_H > 0$ and $Q_H < \bar{Q}_C$ but $c'_H(Q_H) \neq \nu$. [2] use the fact that in equilibrium, when any buyer k deviates by proposing contracts $\{(0, 0), (Q_H + \delta_H, T_H + \epsilon_H)\}$ where $c'_H(Q_H)\delta_H < \epsilon_H < \nu\delta_H$, he should not make a profit. This leads to a contradiction.

⁸In [2], the authors prove that in any equilibrium with $B_H > 0$, marginal rate of substitution for the high type should be equal to ν . In our setting, this translates to $c'_H(Q_H) = \nu$.

The above arguments are valid as long as $Q_H \in (0, \bar{Q}_C)$. However, when $Q_H = \bar{Q}_C$, the deviation contract defined above is not feasible for positive values of δ_H . Therefore, we are able to construct a contradicting argument only for $c'_H(Q_H) > \nu$. This is why $c'_H(Q_H) \leq \nu$ when $B_H > 0$ and $Q_H = \bar{Q}_C$.

Next, we show that the freelancer can trade (Q_H, T_H) with buyers other than k . To do so, we first show that if $U_H > z_H^{-k}(0, 0)$ for some k , then buyer k has a profitable deviation. Indeed, if the freelancer cannot achieve his equilibrium payoff without buyer k 's offer, then buyer k can deviate to the contract $(Q_H, T_H - \epsilon_H)$ for some positive ϵ_H . Such a contract will attract type H for small enough ϵ_H . If type L is not attracted, then the payoff of buyer i satisfies $m_H(B_H + \epsilon_H) > 0$ for small enough ϵ_H . If type L also trades $(Q_H, T_H - \epsilon_H)$, then buyer k 's payoff can be written as $\nu Q_H - T_H + \epsilon_H > 0$ since $T_H = \nu Q_H$ by Proposition 8. Thus, we must have $U_H = z_H^{-k}(0, 0)$. That is, there exists some aggregate trade (Q^{-k}, T^{-k}) that the freelancer can trade with buyers other than k .

Next, we show that $Q^{-k} \neq Q_H$ leads to a contradiction. The proof for the case $Q_H < \bar{Q}_C$ is due to Lemma 4 of Attar et al. [2]. In this case, we have $c'_H(Q_H) = \nu$. This fact and strict convexity of the cost function c_H imply that $T^{-k} > \nu Q^{-k}$. Then, Attar et al. [2] consider two different contracts (q_i, t_i) for $i \in \{H, L\}$ that satisfy the equality $(q_i, t_i) + (Q^{-k}, T^{-k}) = (Q_i, T_i)$ and lead to the desired contradiction.

When $Q_H = \bar{Q}_C$, we have $c'_H(Q_H) \leq \nu$. If $Q^{-k} \neq Q_H$ the only possible case is $Q^{-k} < Q_H$. Even though the aforementioned deviation contracts are feasible, $T^{-k} > \nu Q^{-k}$ is not immediate when $c'_H(Q_H) \neq \nu$. Below, we show that $T^{-k} > \nu Q^{-k}$ even when $c'_H(Q_H) \leq \nu$. Then, the same profitable deviation argument for the aforementioned deviation contracts holds, leading to a contradiction. Observe that

$$U_H = T_H - c_H(Q_H) = \nu Q_H - c_H(Q_H) = T^{-k} - c_H(Q^{-k}).$$

Thus, $T^{-k} = \nu Q^{-k} + \nu(Q_H - Q^{-k}) + c_H(Q^{-k}) - c_H(Q_H)$. Moreover, as the function c_H

is strictly convex, we have the following:

$$c_H(Q_H) = c_H(Q^{-k}) + \int_{Q^{-k}}^{Q_H} c'_H(x) dx < c_H(Q^{-k}) + c'_H(Q_H)(Q_H - Q^{-k}).$$

Combining with the inequality $c'_H(Q_H) \leq \nu$, we obtain the desired result:

$$\begin{aligned} T^{-k} &= \nu Q^{-i} + \nu(Q_H - Q^{-k}) + c_H(Q^{-k}) - c_H(Q_H), \\ &> \nu Q^{-k} + (\nu - c'_H(Q_H))(Q_H - Q^{-k}) \geq \nu Q^{-k}. \end{aligned} \quad \square$$

In the proof of Lemma 6, we first show that in any equilibrium with $B_H > 0$, one must have $c'_H(Q_H) = \nu$. Otherwise, there exists a contract in the neighborhood of (Q_H, T_H) that is profitable if type L is not attracted. If type L also trades the deviation contract, but $c'_H(Q_H) = \nu$ does not hold, then choosing the neighborhood carefully still pays off. The existence of such profitable deviations in equilibrium is a contradiction. However, unlike Attar et al. [2], the deviation contracts designed for the case where type H 's marginal cost is strictly less than ν are not feasible when $Q_H = \bar{Q}_C$. Hence, in our setting, we might have $c'_H(\bar{Q}_C) \leq \nu$ together with $B_H > 0$ in equilibrium. In the remaining part of the result, we show that buyer k has a profitable deviation if the payoff of the freelancer decreases when buyer k withdraws his contract offers. Therefore, in equilibrium, the freelancer should be able to achieve his equilibrium payoff without relying on buyer k . In other words, there exist aggregate trades with buyers other than k that amount to (Q^{-k}, T^{-k}) and satisfy $u_H(Q^{-k}, T^{-k}) = U_H$. When $c'_H(Q_H) = \nu$ and $Q^{-k} \neq Q_H$, strict convexity of the cost function implies $T^{-k} > \nu Q^{-k}$. In this case, buyer k has a profitable deviation. Hence, in equilibrium, the freelancer should be able to trade (Q_H, T_H) with buyers other than k . This is not immediate when $c'_H(Q_H) < \nu$, which is possible in equilibrium when $Q_H = \bar{Q}_C$ in our setup. In such a case, we show that $T^{-k} > \nu Q^{-k}$ still holds, and hence, the same profitable deviation argument applies.

Next, we show that the buyers cannot make aggregate profits with one type of the

freelancer and make losses with the other as long as the high-type freelancer does not trade at the capacity. That is, there is no cross-subsidization in equilibrium unless the high-type freelancer trades at the capacity.⁹

Proposition 9. *In any equilibrium with $Q_H < \bar{Q}_C$, $B_i = 0$ for each i .*

Proof. Suppose to the contrary that we have an equilibrium with $Q_H < \bar{Q}_C$ and $B_i > 0$ for some i . From Corollary 1, we know that i must be H . Then, any buyer k with $b_H^k > 0$ can deviate to the following set of contracts: $\{(0, 0), (Q_L - Q_H + \delta_L, T_L - T_H + \epsilon_L), (q_H^i, t_H^k + \epsilon_H)\}$ where $c'_L(Q_L + \delta_L)\delta_L < \epsilon_L$ and ϵ_H strictly positive. Note that by Lemma 6, type L can trade (Q_H, T_H) with buyers other than k . Combining with the contract $(Q_L - Q_H + \delta_L, T_L - T_H + \epsilon_L)$, type L can trade $(Q_L + \delta_L, T_L + \epsilon_L)$ after the deviation and strictly increase his payoff because of the choice of δ_L and ϵ_L and since we know from Corollary 1 that $S_L = 0$.¹⁰

Now, as in the proof of Lemma 5, fix arbitrary trades from the other buyers' menu of contracts and assume that they amount to (Q^{-k}, T^{-k}) . In equilibrium, we know that $U_L \geq u_L(Q^{-k} + q_H^k, T^{-k} + t_H^k)$ holds for any such Q^{-k} satisfying $Q^{-k} + q_H^k \leq \bar{Q}_C$. Then, for ϵ_H strictly less than $\epsilon_L - \delta_L c'_L(Q_L + \delta_L)$, type L strictly prefers trading $(Q_L - Q_H + \delta_L, T_L - T_H + \epsilon_L)$ with buyer k to trading $(q_H^k, t_H^k + \epsilon_H)$ after the deviation. That is, for any feasible Q^{-k} :

$$\begin{aligned} u_L(Q_L + \delta_L, T_L + \epsilon_L) &= T_L + \epsilon_L - c_L(Q_L + \delta_L), \\ &> T_L + \epsilon_L - c_L(Q_L) - \delta_L c'_L(Q_L + \delta_L), \\ &> U_L + \epsilon_H \geq u_L(Q^{-k} + q_H^k, T^{-k} + t_H^k + \epsilon_H). \end{aligned}$$

⁹Attar et al. [2] shows that no-cross-subsidization holds in any equilibrium in their setting. As we shall see, this is not true in our setup.

¹⁰When $Q_H = \bar{Q}_C$, we must have $Q_H = Q_L = \bar{Q}_C$, and hence deviation contracts become $\{(0, 0), (\delta_L, \epsilon_L), (q_H^k, t_H^k + \epsilon_H)\}$. As negative trades are not allowed, δ_L must be non-negative. Then, type L cannot trade $(Q_L + \delta_L, T_L + \epsilon_L)$ after the deviation and strictly increase his payoff because $Q_L + \delta_L > \bar{Q}_C$ is infeasible. Hence, we cannot rule out cross-subsidization for the case $Q_H = \bar{Q}_C$.

Note that the first inequality above is due to strict convexity of c_L , *i.e.*,

$$c_L(Q_L + \delta_L) = c_L(Q_L) + \int_{Q_L}^{Q_L + \delta_L} c'_L(x) dx < c_L(Q_L) + \delta_L c'_L(Q_L + \delta_L),$$

holds irrespective of the sign of δ_L . After the deviation, type H strictly prefers trading $(q_H^k, t_H^k + \epsilon_H)$ to $(0, 0)$. If type H ends up trading this contract, then buyer k strictly increases his profits for small enough $\delta_L, \epsilon_L, \epsilon_H$, since $S_L = 0$ and $b_H^k > 0$:

$$\begin{aligned} m_L(\nu_L(Q_L - Q_H + \delta_L) - T_L - T_H - \epsilon_L) + m_H(\nu_H q_H^k - t_H^k - \epsilon_H) \\ = m_L(\nu_L \delta_L - \epsilon_L) + m_H(b_H^k - \epsilon_H) > 0. \end{aligned}$$

Hence, due to the zero-profit result, in equilibrium, type H should also trade $(Q_L - Q_H + \delta_L, T_L - T_H + \epsilon_L)$ after the deviation, and buyer k should not make any profit:

$$\nu(Q_L - Q_H + \delta_L) - (T_L - T_H + \epsilon_L) = \nu(Q_L - Q_H + \delta_L) - \nu_L(Q_L - Q_H) - \epsilon_L \leq 0. \quad (3.1)$$

Letting δ_L, ϵ_L go to zero yields $(\nu - \nu_L)(Q_L - Q_H) \leq 0$. Since $\nu > \nu_L$ and $Q_L \geq Q_H$, we must have $Q_L = Q_H$. Then, inequality (3.1) reduces to $\nu \delta_L - \epsilon_L \leq 0$. That is, $\nu \delta_L \leq \epsilon_L$ must hold for any feasible values of δ_L and ϵ_L , that satisfy $c'_L(Q_L) \delta_L < \epsilon_L$. When negative trades are not allowed, and $Q_L = Q_H$, the deviation contract $(Q_L - Q_H + \delta_L, \nu_L(Q_L - Q_H) + \epsilon_L)$ is feasible only for positive values of δ_L . Then, it must be that $c'_L(Q_L) \geq \nu$, for otherwise, one can find a pair of positive δ_L and ϵ_L values such that $\nu \delta_L \leq \epsilon_L$ is not satisfied. Since we have $c'_H(Q_H) = \nu$ due to Lemma 6, this contradicts the assumption $c'_H(Q) > c'_L(Q)$ for any $Q \in [0, \bar{Q}_C]$. \square

The proof of Proposition 9 shows the following contradiction: If $B_H > 0$ in an equilibrium with $Q_H < \bar{Q}_C$, then some buyer k making profits with the high-type freelancer has a profitable deviation. The following deviation contracts from Attar et al. [2] are useful in our setting: no-trade contract, a contract for the additional aggregate trade made with the

low-type, $(Q_L - Q_H, T_L - T_H)$, and his equilibrium trade with the high-type, (q_H^k, t_H^k) . In the proof, the last two contracts are slightly altered so that the low-type freelancer prefers trading $(Q_L - Q_H, T_L - T_H)$ after the deviation, whereas the high-type prefers trading either of the last two contracts. Note that due to Lemma 6, when $B_H > 0$, the high-type freelancer can achieve his aggregate equilibrium trades, (Q_H, T_H) , without relying on buyer k 's offer. Hence, after the deviation, the low-type freelancer can trade $(Q_L - Q_H, T_L - T_H)$ with buyer k and trade (Q_H, T_H) with the other buyers. Since $S_L = 0$ by Corollary 1, after the deviation, buyer k makes zero profit with the low-type freelancer. On the other hand, if the high-type freelancer trades (q_H^k, t_H^k) after the deviation, then buyer k makes a strictly positive profit. Hence, the high-type freelancer should also trade $(Q_L - Q_H, T_L - T_H)$ after the deviation, and the expected profit of buyer k should be at most zero. This is only possible if $Q_L = Q_H$ and $c'_L(Q_L) \geq \nu$, since otherwise, there exists a contract in the neighborhood of $(Q_L - Q_H, T_L - T_H)$ that leads to a profitable deviation. Since $c'_H(Q_H) = \nu$ when $Q_H < \bar{Q}_C$ from Lemma 6, this contradicts the assumption $c'_H(Q) > c'_L(Q)$ for any $Q \in [0, \bar{Q}_C]$.

The deviation contracts used in the proof of Proposition 9 are not feasible when $Q_H = \bar{Q}_C$, since they require a service level that is greater than \bar{Q}_C . This is why we require Q_H to be strictly less than the capacity in the statement of Proposition 9.

Proposition 10 below puts more detail on any equilibrium with no-cross-subsidization: We understand that in such an equilibrium, each traded contract yields zero profit, and type H chooses not to trade.

Proposition 10. *In any equilibrium with $B_H = B_L = 0$, $b_i^k = 0$ and $q_L^k \geq q_H^k = 0$ for all i and k .*

Proof. Consider an equilibrium with $B_H = B_L = 0$. By definition, $B_H = 0$ implies $T_H = \nu_H Q_H$. On the other hand, we have $T_H = \nu Q_H$ due to Proposition 8. This is only possible if $(Q_H, T_H) = (0, 0)$. Since only non-negative trades are allowed, we also have $q_H^k = 0$ for all k . If $t_H^k = 0$ for all k , then the proof is complete due to the zero-profit

result. Otherwise, there exists some buyer k with $t_H^k < 0$. In this case, type H can strictly increase his profit by trading $(0, 0)$ with buyer k , a contradiction. \square

Propositions 9 and 10 lead to the following immediate result: In any equilibrium, the high-type freelancer either does not trade or his capacity constraint is binding, which we formalize below.

Corollary 2. *In any equilibrium, either $Q_H = 0$ or $Q_H = \bar{Q}_C$.*

Lemma 7 below derives necessary conditions regarding the marginal cost of both types of the freelancer for any equilibrium with no-cross-subsidization.

Lemma 7. *In any equilibrium with $B_H = B_L = 0$, if $Q_L > 0$, then*

$$c'_L(Q_L) \begin{cases} = \nu_L & \text{if } Q_L < \bar{Q}_C, \\ \leq \nu_L & \text{if } Q_L = \bar{Q}_C. \end{cases}$$

Moreover, if $Q_i = 0$, then $c'_i(0) \geq \min \{\nu_i, \nu\}$ for $i \in \{H, L\}$.

Proof. For the case $Q_L > 0$ and $Q_L < \bar{Q}_C$, the result $c'_L(Q_L) = \nu_L$ follows from Lemma 6 of Attar et al. [2]. Suppose that in an equilibrium with $B_H = B_L = 0$, we have $Q_L \in (0, \bar{Q}_C)$ and $c'_L(Q_L) \neq \nu_L$. Then, any buyer k can propose the contract $(Q_L + \delta_L, T_L + \epsilon_L)$ for some small δ_L and ϵ_L that satisfy $c'_L(Q_L)\delta_L < \epsilon_L < \nu_L\delta_L$. Attar et al. [2] show that this is a profitable deviation, contradicting the fact that we are in an equilibrium.

On the other hand, when $Q_L = \bar{Q}_C$, the deviation contract defined above is infeasible for the positive values of δ_L and ϵ_L . But, $c'_L(Q_L)\delta_L < \epsilon_L < \nu_L\delta_L$ can still be satisfied for negative values of δ_L and ϵ_L . Therefore, a similar contradiction arises only for $c'_L(Q_L) > \nu_L$ in this case. Hence, $c'_L(Q_L) \leq \nu_L$ when $Q_L = \bar{Q}_C$.

Finally, the result $c'_i(0) \geq \min \{\nu_i, \nu\}$ whenever $Q_i = 0$ follows from Lemma 5 of [2]. \square

Lemma 7 presents conditions for equilibria with no-cross-subsidization for the cases $Q_L > 0$, $Q_L = 0$, and $Q_H = 0$. We do not consider the case $Q_H > 0$ since we know that the high-type freelancer does not trade in any equilibrium with no-cross-subsidization by Proposition 10.¹¹

3.4 The Main Results

The results that we obtain in the previous section lead us to our first main result, which provides a characterization of aggregate equilibrium trades as well as necessary conditions for the existence of a pure-strategy equilibrium:

Theorem 7. *If an equilibrium exists, then $\nu \leq c'_H(0)$ or $c'_H(\bar{Q}_C) \leq \nu$. Moreover, the following statements hold.*

- (i) *If $c'_H(\bar{Q}_C) \leq \nu$, all equilibria are pooling with $Q_L = Q_H = \bar{Q}_C$.*
- (ii) *If $\nu_L \leq c'_L(0)$ and $\nu \leq c'_H(0)$, all equilibria are pooling with $Q_L = Q_H = 0$.*
- (iii) *If $c'_L(0) < \nu_L$ and $\nu \leq c'_H(0)$, all equilibria are separating with:*

$$Q_L = \begin{cases} Q_L^* & \text{if } c'_L(\bar{Q}_C) > \nu_L, \\ \bar{Q}_C & \text{if } c'_L(\bar{Q}_C) \leq \nu_L, \end{cases} \text{ and } Q_H = 0,$$

where Q_L^* satisfies $c'_L(Q_L^*) = \nu_L$.

Proof. Suppose that an equilibrium exists. First, observe that the hypothesis of (i), (ii), and (iii) are mutually exclusive and collectively exhaustive for $\nu \leq c'_H(0)$ or $c'_H(\bar{Q}_C) \leq \nu$.

¹¹In the proof of Lemma 7, we use the deviation contracts that are presented in a similar result by Attar et al. [2], but we adjust the arguments according to the feasibility of these contracts.

Furthermore, by Corollary 2, in any equilibrium, either $Q_H = 0$ or $Q_H = \bar{Q}_C$. Therefore, the conclusions of the implications of (i), (ii), and (iii) are mutually exclusive and collectively exhaustive for all possible aggregate equilibrium trades.

- (i) If $Q_H = \bar{Q}_C$, then we must have a pooling equilibrium with $Q_L = Q_H = \bar{Q}_C$, since we know that $\bar{Q}_C \geq Q_L \geq Q_H$ holds in any equilibrium. In this case, $B_H > 0$ follows from Proposition 8, and then $c'_H(\bar{Q}_C) \leq \nu$ follows from Lemma 6.

On the other hand, if $Q_H < \bar{Q}_C$, then two different equilibria may occur. In either of these, there is no cross-subsidization due to Proposition 9, and $Q_L \geq Q_H = 0$ by Proposition 10.

- (ii) If it is a pooling equilibrium, then it must be $Q_L = Q_H = 0$. In this case, we have $c'_i(0) \geq \min \{\nu_i, \nu\}$ for $i \in \{H, L\}$ due to Lemma 7.
- (iii) If it is a separating equilibrium, then it must satisfy $Q_L > Q_H = 0$. Depending on the feasibility of Q_L^* , two different cases are possible for Q_L : If $Q_L^* < \bar{Q}_C$, then by the strict convexity of c_L , we have $c'_L(\bar{Q}_C) > \nu_L$. In this case, one must have $Q_L = Q_L^*$ by Lemma 7. If $Q_L^* \geq \bar{Q}_C$, then by the strict convexity of c_L , we have $c'_L(\bar{Q}_C) \leq \nu_L$. Then by Lemma 7, we have $Q_L = \bar{Q}_C$. In either case, $c'_L(0) < \nu_L$ follows from strict convexity of c_L , and $c'_H(0) \geq \nu$ follows from Lemma 7.

Finally, notice that the cases above together imply that an equilibrium exists only if $\nu \leq c'_H(0)$ or $c'_H(\bar{Q}_C) \leq \nu$ holds. Since both the hypotheses and the conclusions of (i), (ii), and (iii) are mutually exclusive and collectively exhaustive of all aggregate equilibrium trades, the proof is complete. \square

The proof of Theorem 7 characterizes the necessary conditions both for the pooling and the separating equilibria. If there exists an equilibrium with $Q_H < \bar{Q}_C$, then by Proposition 9, we know that there is no cross-subsidization. In this case, the aggregate equilibrium trades must satisfy $Q_L \geq Q_H = 0$ by Proposition 10. Hence, in a pooling equilibrium with

no-cross-subsidization, both types of the freelancer must not trade. Lemma 7 gives the no-trade-equilibrium conditions on the marginal costs of both types as $c'_i(0) \geq \min \{\nu_i, \nu\}$ for $i \in \{H, L\}$. On the other hand, in a separating equilibrium with no-cross-subsidization, we must have $Q_L > Q_H = 0$. In this case, the aggregate equilibrium trade of the low-type freelancer depends on the feasibility of Q_L^* . If it is feasible, then the low-type freelancer trades efficiently in equilibrium. Otherwise, he will trade at the capacity. In either of the cases, strict convexity of the cost function implies $c'_L(0) < \nu_L$, whereas $c'_H(0) \geq \nu$ follows from Lemma 7.

Since $\bar{Q}_C \geq Q_L \geq Q_H$ in any equilibrium, it follows from Corollary 2 that the only remaining case is a pooling equilibrium where both types trade at the capacity. Due to Proposition 8, the aggregate equilibrium trades, in this case, are characterized by $Q_H = Q_L = \bar{Q}_C$ and $T_H = T_L = \nu \bar{Q}_C$, which result in cross-subsidization. Then by Proposition 8, the buyers make aggregate profits with the high-type freelancer. Hence, by Lemma 6, in such an equilibrium, we have $c'_H(\bar{Q}_C) \leq \nu$.

Next, we show that any aggregate equilibrium trade can be supported by at least two buyers posting the same linear tariffs (described in Theorem 8). Furthermore, the necessary conditions for equilibrium existence given in Theorem 7 are also sufficient.

Theorem 8. *An equilibrium exists if and only if $\nu \leq c'_H(0)$ or $c'_H(\bar{Q}_C) \leq \nu$. Moreover, the following statements hold.*

- (i) *If $\nu \leq c'_H(0)$, any equilibrium can be supported by at least two buyers posting the same tariff*

$$t(q) = \nu_L q, \quad 0 \leq q \leq \bar{Q}_C,$$

while the other buyers remain inactive.

- (ii) *If $c'_H(\bar{Q}_C) \leq \nu$, any equilibrium can be supported by at least two buyers posting the same tariff*

$$t(q) = \nu q, \quad 0 \leq q \leq \bar{Q}_C,$$

while the other buyers remain inactive.

Proof. We first show that if $\nu \leq c'_H(0)$, then there exists an equilibrium. Fix an integer K satisfying $2 \leq K \leq n$ and suppose that K buyers post the tariff given in (i) while the other buyers remain inactive. This means, in the aggregate, competitors of any buyer post the tariff $T^-(Q^-) = \nu_L Q^-$ for $0 \leq Q^- \leq \bar{Q}$ where \bar{Q} is either $K\bar{Q}_C$ or $(K-1)\bar{Q}_C$. Note that \bar{Q} cannot be smaller than \bar{Q}_C since $K \geq 2$. Suppose a buyer deviates and ends up trading the contracts (q_L, t_L) and (q_H, t_H) with types L and H , respectively. At least one of these contracts should give positive profits if the deviating buyer has a profitable deviation.

First, we consider the contract (q_H, t_H) : For this contract to give positive profit, we must have $\nu_H q_H > t_H$. When only non-negative trades are allowed, we can directly deduce that $q_H \in (0, \bar{Q}_C]$. Define $Q_i^- \in [0, \bar{Q}_C]$ as the quantity type $i \in \{H, L\}$ trades with the deviator's competitors after the deviation. Define the total quantity traded by $i \in \{H, L\}$ as $\hat{Q}_i = q_i + Q_i^-$ and similarly $\hat{T}_i = t_i + T^-(Q_i^-)$. Since type H prefers trading (\hat{Q}_H, \hat{T}_H) , we have:

$$u_H(\hat{Q}_H, \hat{T}_H) \geq u_H(0, 0).$$

Together with $\nu \leq c'_H(0)$, the above inequality implies $\hat{T}_H > \nu \hat{Q}_H$ since otherwise, type H would prefer no-trade to (\hat{Q}_H, \hat{T}_H) .

Let us consider the payoff of type L if he also trades (q_H, t_H) with the deviator. He would choose a feasible Q^- maximizing $u_L(q_H + Q^-, t_H + T^-(Q^-))$. This optimization problem is subject to the following feasibility constraints: $0 \leq Q^- \leq \bar{Q}_C - q_H$. We now show that $Q^- = \hat{Q}_L - q_H$ is a feasible solution. Firstly, the capacity constraint is satisfied since $\hat{Q}_L \leq \bar{Q}_C$ by definition. For the non-negativity constraint, recall that $\hat{Q}_L \geq \hat{Q}_H$ holds due to the assumption $c'_H(Q) > c'_L(Q)$ for any $Q \in [0, \bar{Q}_C]$. Then, we have $\hat{Q}_L - q_H \geq Q_H^- \geq 0$. Thus, after the deviation, type L can receive at least $u_L(\hat{Q}_L, t_H + T^-(\hat{Q}_L - q_H))$. By the definition of the tariff given in (i), we can rewrite the

aggregate transfers as follows:

$$t_H + T^-(\hat{Q}_L - q_H) = \hat{T}_H + T^-(\hat{Q}_L - q_H) - T^-(Q_H^-) = \hat{T}_H + \nu_L(\hat{Q}_L - \hat{Q}_H).$$

Since type L prefers trading (\hat{Q}_L, \hat{T}_L) , it follows that $\hat{T}_L \geq \hat{T}_H + \nu_L(\hat{Q}_L - \hat{Q}_H)$. Then, the aggregate profit written below can be at most zero because $\hat{T}_H > \nu\hat{Q}_H$:

$$\nu\hat{Q}_H - \hat{T}_H + m_L[\nu_L(\hat{Q}_L - \hat{Q}_H) - (\hat{T}_L - \hat{T}_H)] \leq 0.$$

By the definition of the tariff given in (i), the competitors of the deviator cannot make losses. Hence, the deviator does not have a profitable deviation.

Next, we consider the contract (q_L, t_L) : For this contract to give a positive profit, we must have $\nu_L q_L > t_L$, and q_L must be strictly positive. By the definition of the tariff given in (i), competitors of the deviator cannot make losses. Hence, $\nu_L q_L > t_L$ implies $\nu_L \hat{Q}_L > \hat{T}_L$. Thus,

$$u_L(\hat{Q}_L, \nu_L \hat{Q}_L) > u_L(\hat{Q}_L, \hat{T}_L).$$

We know that the inequalities $\hat{Q}_L \leq \bar{Q}_C \leq \bar{Q}$ hold. Since type L can trade $(\hat{Q}_L, \nu_L \hat{Q}_L)$ with the competitors of the deviator, we arrive at a contradiction.

Now, keeping the same notation introduced above, we assume that $c'_H(\bar{Q}_C) \leq \nu$ and let $K \geq 2$ many buyers offer the tariff given in (ii) while the others remain inactive. In this case, competitors of any buyer post $T^-(Q^-) = \nu Q^-$ in aggregate for $0 \leq Q^- \leq \bar{Q}$ where \bar{Q} is either $K\bar{Q}_C$ or $(K-1)\bar{Q}_C$. Suppose a buyer deviates and ends up trading the contracts (q_L, t_L) , and (q_H, t_H) with the types L and H , respectively. Defining $\hat{Q}_i, \hat{T}_i, Q_i^-$

and T_i^- as before, we must have:

$$\begin{aligned}
u_H(\hat{Q}_H, \hat{T}_H) &\geq u_H(q_H + (\bar{Q}_C - q_H), t_H + T_H^-(\bar{Q}_C - q_H)), \\
t_H + T_H^-(Q_H^-) - c_H(\hat{Q}_H) &\geq t_H + T_H^-(\bar{Q}_C - q_H) - c_H(\bar{Q}_C), \\
c_H(\bar{Q}_C) - c_H(\hat{Q}_H) &\geq T_H^-(\bar{Q}_C - \hat{Q}_H), \\
\int_{\hat{Q}_H}^{\bar{Q}_C} c'_H(x) dx &\geq \nu(\bar{Q}_C - \hat{Q}_H),
\end{aligned}$$

since type H can trade $\bar{Q}_C - q_H$ with the competitors of the deviator who offer the tariff given in (ii). Since c_H is strictly convex and $c'_H(\bar{Q}_C) \leq \nu$, the inequalities above can be satisfied only if $\hat{Q}_H = \bar{Q}_C$. Then $\hat{Q}_L \geq \hat{Q}_H$ implies that both types trade at the same aggregate level and they are indifferent between trading $(q_H, t_H) + (Q_H^-, T_H^-)$ and $(q_L, t_L) + (Q_L^-, T_L^-)$. Since the deviator does not know the type of the freelancer, his expected payoff from the contract (q_i, t_i) , if traded by the freelancer, can be written as $\nu q_i - t_i$ for $i \in \{H, L\}$. On the other hand, if the freelancer prefers trading (q_i, t_i) over the contracts of the tariff given in (ii), then it must be that $t_i \geq \nu q_i$ for $i \in \{H, L\}$. Hence, the deviator does not have a profitable deviation. \square

The proof of Theorem 8 shows that when either of the necessary conditions presented in Theorem 7 is satisfied, then at least two buyers offering the corresponding tariff leads to an equilibrium. Therefore, either of the necessary conditions presented in Theorem 7 is also sufficient for the existence of a pure-strategy equilibrium.

To sum up, our main results fully characterize the aggregate equilibrium trades of a freelancer under nonexclusive competition when there is adverse selection.¹² We provide

¹²Our proofs employ deviation contracts tailored to each type of the freelancer, similar in spirit to Attar et al. [2], who note that “*beyond two types, it becomes hard, if not intractable, to control the behavior of each type following such a deviation*” under nonexclusive competition. In a subsequent paper [31], they extend their model to arbitrary discrete distributions under the assumption that the buyers offer convex menus, ensuring that the seller’s indirect utility functions satisfy a single-crossing property. In the present chapter, we do not impose any restrictions on the contract offers, and hence, unable to speculate beyond two-type settings.

necessary and sufficient conditions for the equilibrium existence and show that if an equilibrium exists, then the aggregate equilibrium trades are unique. In equilibrium, each buyer makes zero profit in expectation, even though they can make a positive profit from a contract traded with the high-type freelancer. Furthermore, any equilibrium can be supported by linear tariffs. Depending on the preferences of the high-type freelancer, details of equilibria can be summarized as follows. If the high-type freelancer is not willing to serve at a price equal to the expected quality, then we obtain an Akerlof-like result: The high-type freelancer does not trade in equilibrium, and there is no cross-subsidization. In this case, the cost function of the low-type freelancer determines his aggregate equilibrium trades: He might trade efficiently, not trade at all, or exhaust her capacity. On the other hand, if the high-type freelancer is willing to serve for the price equal to the expected quality at every feasible level, then both types exhaust their capacity in every equilibrium, and there is cross-subsidization.

3.5 Concluding Remarks

Our results point out an Akerlof-like market breakdown as the high-type freelancer prefers not to trade in equilibrium—when he is not willing to trade at a unit price equal to the expected quality of the service. One implication of this result is that if the high-type freelancer can credibly signal his type to the buyers, then he may increase his payoff by extracting rents from the buyers due to nonexclusive competition. Hence, it is reasonable to expect the rise of intermediaries that help freelancers to signal their types credibly by exerting a cost. This is indeed the case for many labor markets: There are intermediaries for freelancers such as Upwork (formerly ODesk), peopleperhour.com, or guru.com. As noted by Pelletier & Thomas [32], “*missing information hampers the level of activity in these markets.*” In other words, these markets are negatively affected by adverse selection. Hence, according to our results, prices observed in these markets should be equal to the expected quality of the service. Although these platforms provide information on

the freelancer to hinder the effects of adverse selection, the freelancers who are new to the market do not have much to offer in this aspect. As reported by Pallais [33], some of the entry-level freelancers of ODesk are “*inefficiently unemployed*” due to uncertainty about their abilities. Building on this result, Stanton & Thomas [34] note the emergence of intermediaries that enables the freelancers to signal their quality. Their results suggest that the freelancers who are not affiliated with any of these intermediaries earn substantially less at the beginning of their careers compared to similar freelancers who have an affiliation. In other words, the buyers protect themselves against adverse selection by offering low prices to the entry-level freelancers. Therefore, these empirical results are in line with our theoretical findings.

There are minor differences between our setting and those of Attar et al. [1] and Attar et al. [2]. One question that comes to mind is whether it is possible to find a direct relationship between the changes in the problem settings and the differences in the results. In [2], the authors compare their results to those of [1] and conclude that no-cross-subsidization result is not attainable in [1] because of the capacity constraint. Our results confirm this assessment. When the marginal cost of the high-type freelancer is smaller than the expected quality at any feasible service level, there exists a pooling equilibrium with cross-subsidization. On the other hand, when the marginal cost of the high-type freelancer is greater than the expected quality at the no-trade point, all equilibria exhibit no-cross-subsidization. Hence, the aggregate trades derived in our setting resemble those of [1], but in our setup, equilibrium does not need to exist. As in [2], preferences of the high-type freelancer regulate the existence of a pure-strategy equilibrium while preferences of the low-type freelancer shape the aggregate equilibrium trades.

Finally, we highlight the welfare implications of the capacity constraint, which is the key difference between our study and Attar et al. [2]. For the case of non-negative trades, Attar et al. [2] find that the high-type freelancer remains inactive in any pure-strategy equilibrium. Our results suggest that adding a capacity constraint to the freelancer’s side does not disturb these equilibria, but the capacity constraint may prevent the low-type

freelancer from trading efficiently. Hence, he may be worse off because he is unable to realize some of the profits due to the capacity constraint. On the other hand, the capacity constraint leads to an additional pooling equilibrium. Both types exhaust their capacity at a price equal to the expected quality in this equilibrium, whereas the buyers continue to make zero profit. We see that the low-type freelancer is better off while the high-type freelancer is worse off in this pooling equilibrium when compared to the complete information case with the capacity constraint.

Chapter 4

Prophet Inequalities & Polyhedrons

4.1 Introduction

Consider a gambler who wants to maximize his gains from a box opening game. The rules of the game are as follows. Each box contains a non-negative reward, independent draws from known distributions (not necessarily identical). The contents of the boxes are revealed sequentially to the gambler as the game proceeds. Each time a box is opened, the gambler must choose one of the following. He can either claim the last unboxed reward or discard it and move on. How should the gambler play this game to maximize his expected gain? Optimal stopping literature, as its name suggests, deals with the extensions of this problem, while we will be interested in another dimension. Note that the above-mentioned problem can be solved via Dynamic programming.

Now, consider a prophet who can foresee all reward realizations to come. Such a player can choose the highest reward in any instance of the game, hence gaining the expected value of the maximum reward. Krengel & Sucheston [35] showed that the optimal strategy

of the gambler yields at least half of the optimal strategy of the prophet. This approximation result has become known as the prophet inequality, and it is known to be tight, *i.e.*, there is an instance where the gambler cannot gain more than half of what the prophet gains. This seminal work inspired many others to study prophet inequalities in different fields, including Mechanism design.

In this chapter, we use tools from linear programming to obtain prophet inequalities for extensions of the gambler’s problem where he is constrained to choose a strategy from a given polyhedron. This methodology was first used by Epiropou & Vohra [36], who also inspired this work. Epiropou & Vohra [36] consider the *online* version of allocation with costly verification problem (see the introduction of Chapter 2 for the offline version). In their words, “*the agents arrive and depart one at a time, and the decision to allocate the object to an agent must be made upon the arrival of an agent. If the principal declines to allocate the object to an agent, the agent departs and cannot be recalled. If the principal allocates the object to an agent, the decision is irreversible.*” They showed that this problem could be posed as a compact linear program, which allows them to derive a prophet inequality through feasibility arguments. Similarly, the main idea of our methodology is to find a feasible strategy for the gambler that chooses each box at least half as much as the prophet in expectation. Compared to the prophet inequality literature methods, such as dynamic programming, induction, conditioning, conjugate duality, algebraic inequalities, and moment theory (see Hill & Kertz [37] for details), this approach gives a systematic and straightforward way of deriving prophet inequalities.

We contribute to the literature by re-establishing three different results. Note that we assume independently drawn rewards unless otherwise stated.

- (i) $\frac{1}{2}$ -prophet inequality when the player is restricted to a polymatroid, which has non-negative coefficients in its unique Minkowski sum of simplices.
- (ii) $\frac{1}{n}$ -prophet inequality when there are n boxes, the reward distributions can be dependent, and the gambler is restricted to an arbitrary polyhedron.

(iii) $\frac{1}{K+1}$ -prophet inequality when the gambler is subject to K many constraints.

Although our contribution is solely methodological for now, we hope that the simplicity of our method may lead to novel prophet inequalities.

Literature Review. The literature on prophet inequalities is vast. An interested reader can see Hill & Kertz [37] for a classic survey. Here, we intend to review the recent work that links prophet inequalities to robust mechanism design. Mechanism design literature traditionally focuses on the offline setting where the set of buyers is static, and each agent is present when the mechanism undergoes. This assumption is not fit for online markets where buyers arrive and depart dynamically. Hence, parallel to the rise of online auctions, literature on online mechanism design has accumulated since 2000. The first paper that utilizes prophet inequalities to obtain a mechanism with a constant-approximation guarantee is by Hajiaghayi et al. [38], who leveraged the monotonicity property of some algorithms from the prophet inequality literature to find direct online mechanisms. Then, Chawla et al. [39] showed that one could guarantee to obtain half of the payoff of the optimal offline Bayesian mechanism using sequential posted price mechanisms, which are especially important for the online setting as they are easy to use. Alaei [40] improved the prophet inequality for the setting where the gambler can choose k many random variables before stopping and applied it to a more general mechanism design framework than [39]. Finally, Kleinberg & Weinberg [41] considered the setting where the gambler is subject to a matroid constraint and proved the $\frac{1}{2}$ -prophet inequality by constructing an algorithm for the gambler, whereas Dütting & Kleinberg [42] extended this result to polymatroids. In [42], the authors showed that polymatroids could be reduced to block-structured matroids and utilized the algorithm developed by [41]. Both [41] and [42] discuss the applications of their results to mechanism design problems. For further reading on the literature that connects prophet inequalities to mechanism design, one can refer to Lucier [43].

4.2 Model

The gambler must select an element (called a ‘box’) with an associated reward in the original problem. There are n boxes, and r_i denotes the reward of the i^{th} box, which is a random draw from a known distribution F_i (with density f_i). We will deal with a *general* version of the fractional prophet problem, in which the gambler is allowed to acquire a fraction of a reward. The generality comes from the fact that the gambler can choose any subsets of boxes from a prescribed family of subsets, which a collection of inequalities can describe.

We first define the variables for the online problem. Let $q_i(r_1, \dots, r_i)$ be the fraction of reward acquired from box $i = 1, \dots, n$, given that the profile of rewards (r_1, \dots, r_i) was realized. As in the mechanism design literature, we will refer to q as ex-post allocation variables. For all i , let

$$\vec{q}(r_1, \dots, r_i) = [q_1(r_1), \dots, q_i(r_1, \dots, r_i)].$$

To describe the set of feasible ex-post allocations, let A be a non-negative $K \times n$ matrix. Let A_i be the submatrix of A containing the first i columns of A and denote the j^{th} column of A by a^j . Then, \vec{q} is feasible if for all $1 \leq i \leq n$

$$A_i \vec{q}(r_1, \dots, r_i) \leq b \quad \forall r_1, \dots, r_i.$$

It will sometimes be helpful to write this expression down completely. If a_{kj} denotes the entry in the k^{th} row and j^{th} column, we require that for all $i \leq n$, for all $k \in \{1, \dots, K\}$ and all profiles (r_1, \dots, r_i) :

$$\sum_{j \leq i} a_{kj} q_j(r_1, \dots, r_j) \leq b_k.$$

We describe a more compact formulation using the variables $Q_j(r_j)$ defined as follows:

$$Q_j(r_j) = \mathbb{E}_{r_1, \dots, r_{j-1}}[q_j(r_1, \dots, r_j)].$$

The variables $Q_j(r_j)$ are sometimes called interim allocations. A formulation of the prophet problem in terms of the interim allocations is called the reduced form representation.

An interim allocation Q is implementable up until round i if there exists an ex-post allocation \vec{q} satisfying the following set of constraints ($FP^i[Q]$):

$$\begin{aligned} A_i \vec{q}(r_1, \dots, r_i) &\leq b && \forall r_1, \dots, r_i, \\ \mathbb{E}_{r_1, \dots, r_{j-1}}[q_j(r_1, \dots, r_j)] &\geq Q_j(r_j) && \forall r_j, \forall j \leq i, \\ \vec{q}(r_1, \dots, r_i) &\geq 0 && \forall r_1, \dots, r_i. \end{aligned}$$

If the interim allocation Q is implementable up until round i , we will say that $FP^i[Q] \neq \emptyset$. An interim allocation Q is implementable if $FP^i[Q] \neq \emptyset$ for all $i = 1, \dots, n$.

Consider the following optimization problem:

$$\begin{aligned} h_{i+1}(Q_1, \dots, Q_i) &= \max_{q_1, \dots, q_i} \mathbb{E}_{r_1, \dots, r_i}[z(r_1, \dots, r_i)] \\ \text{s.t. } A_i \vec{q}(r_1, \dots, r_i) + a^{i+1} z(r_1, \dots, r_i) &\leq b \quad \forall r_1, \dots, r_i, \end{aligned} \quad (4.1)$$

$$\mathbb{E}_{r_1, \dots, r_{j-1}}[q_j(r_1, \dots, r_j)] \geq Q_j(r_j) \quad \forall r_j, \forall j \leq i, \quad (4.2)$$

$$\vec{q}, z \geq 0. \quad (4.3)$$

Lemma 8. Q is implementable if and only if $Q_{i+1}(r_{i+1}) \leq h_{i+1}(Q_1, \dots, Q_i)$ for all rewards r_{i+1} and all $i = 1 \dots, n - 1$.

Proof. If Q is implementable, the statement is clearly true. So, suppose to the contrary that $Q_{i+1}(r) \leq h_{i+1}(Q_1, \dots, Q_i)$ for all rewards r and all $i = 1 \dots, n$, but Q is

not implementable, *i.e.*, $FP^{i+1}[Q] = \emptyset$. Let (\vec{q}, z) be an optimal solution that yields $h_{i+1}(Q_1, \dots, Q_i)$ and therefore satisfies (4.1-4.3). Notice, \vec{q} is a feasible ex-post allocation in $FP^i[Q]$. Then, there does not exist any q_{i+1} that solves the following system of inequalities together with \vec{q} :

$$A_i \vec{q}(r_1, \dots, r_i) + a^{i+1} q_{i+1}(r_1, \dots, r_{i+1}) \leq b \quad \forall r_1, \dots, r_{i+1}, \quad (4.4)$$

$$\mathbb{E}_{r_1, \dots, r_i} [q_{i+1}(r_1, \dots, r_{i+1})] \geq Q_{i+1}(r_{i+1}) \quad \forall r_{i+1}, \quad (4.5)$$

$$\vec{q}, z \geq 0. \quad (4.6)$$

Let y_{r_1, \dots, r_i} be the optimal dual variables associated with (4.1). Then, for all (r_1, \dots, r_{i+1}) :

$$\begin{aligned} y_{r_1, \dots, r_i} A_i \vec{q}(r_1, \dots, r_i) + y_{r_1, \dots, r_i} a^{i+1} q_{i+1}(r_1, \dots, r_{i+1}) &\leq y_{r_1, \dots, r_i} b, \\ q_{i+1}(r_1, \dots, r_{i+1}) &\leq \frac{y_{r_1, \dots, r_i} b - y_{r_1, \dots, r_i} A_i \vec{q}(r_1, \dots, r_i)}{y_{r_1, \dots, r_i} a^{i+1}}. \end{aligned}$$

We can use this to upper bound each $q_{i+1}(r_1, \dots, r_{i+1})$. Knowing that (4.4-4.6) is infeasible implies that even if we set each $q_{i+1}(r_1, \dots, r_{i+1})$ at its upper bound, we must violate one of the constraints in (4.5). Hence, there is an r_{i+1} such that:

$$\begin{aligned} Q_{i+1}(r_{i+1}) &> \mathbb{E}_{r_1, \dots, r_i} \left[\frac{y_{r_1, \dots, r_i} b - y_{r_1, \dots, r_i} A_i \vec{q}(r_1, \dots, r_i)}{y_{r_1, \dots, r_i} a^{i+1}} \right] \\ &= \mathbb{E}_{r_1, \dots, r_i} [z(r_1, \dots, r_i)] = h_{i+1}(Q_1, \dots, Q_i). \end{aligned}$$

The penultimate equation follows by complementary slackness, and we obtain a contradiction. \square

Now we describe a linear programming formulation for the problem solved in hindsight, *i.e.*, the offline version. In the offline problem, the fraction of reward acquired from i is based on the *entire* profile of rewards. Denote by $w_i(r_1, \dots, r_n)$, the ex-post allocation to

i at reward profile (r_1, \dots, r_n) . Let $\vec{w}(r_1, \dots, r_n) = [w_1(r_1, \dots, r_n), \dots, w_n(r_1, \dots, r_n)]$. As before, we define the interim allocation probabilities as follows:

$$W_i(r_i) = \mathbb{E}_{r_{-i}}[w_i(r_i, r_{-i})].$$

Here r_{-i} denotes the profile $(r_1, \dots, r_{i-1}, r_{i+1}, \dots, r_n)$. The offline problem can be expressed as follows:

$$\begin{aligned} \max \quad & \sum_{j=1}^n \mathbb{E}_{r_j}[r_j W_j(r_j)] \\ \text{s.t.} \quad & A_i \vec{w}(r_1, \dots, r_n) \leq b && \forall r_1, \dots, r_n, \forall i \leq n, \\ & W_i(r_i) = \mathbb{E}_{r_{-i}}[w_i(r_i, r_{-i})] && \forall r_i, \forall i, \\ & \vec{w}(r_1, \dots, r_n) \geq 0 && \forall r_1, \dots, r_n. \end{aligned}$$

Denote the optimal solution to the offline problem by (\vec{w}^*, W^*) .

If we can show that $h_{i+1}(0.5W_1^*, \dots, 0.5W_i^*) \geq 0.5W_{i+1}^*(r_{i+1})$ for all r_{i+1} , then, from Lemma 8, it follows that an online solution with a value of at least half the optimal offline solution is implementable. In other words, we obtain a prophet inequality with a bound of $1/2$. We first apply this idea to the case when $K = 1$, *i.e.*, there is only one constraint. Note that Epitropou & Vohra [36] are the first to consider this methodology in a mechanism design setting with one constraint. They prove the feasibility of $0.5W^* = (0.5W_1^*, \dots, 0.5W_n^*)$ in their online problem to obtain a $\frac{1}{2}$ -prophet inequality. We present an alternative proof for a more general class of singly constrained stopping problems.

Proposition 11. *Given a scalar b and a non-negative n dimensional vector A , if the gambler is subject to the constraints $A\vec{q}(r_1, \dots, r_n) \leq b$ for all r_1, \dots, r_n , then the interim allocation $0.5W^*$ is implementable.*

Proof. Notice that $(0.5W_1^*, 0.5W_2^*)$ is feasible in any online problem with nonnegative A

thanks to scaling by 0.5. Hence, we can proceed as follows. Assume that there exists some $i < n$ such that $h_j(0.5W_1^*, \dots, 0.5W_{j-1}^*) \geq 0.5W_j^*(r_j)$ for all r_j and $j \leq i$. Then, the problem $h_{i+1}(0.5W_1^*, \dots, 0.5W_i^*)$ is feasible. Therefore, for any feasible \vec{q} with the expected values $(0.5W_1^*, \dots, 0.5W_i^*)$, it is optimal to set:

$$z^*(r_1, \dots, r_i) = b - \sum_{j \leq i} a_j q_j(r_1, \dots, r_j) \quad \forall r_1, \dots, r_i,$$

as there is only one constraint for each r_1, \dots, r_i . Then, the objective becomes:

$$\begin{aligned} h_{i+1}(0.5W_1^*, \dots, 0.5W_i^*) &= \mathbb{E}_{r_1, \dots, r_i}[z^*(r_1, \dots, r_i)] = \mathbb{E}_{r_1, \dots, r_i}[b - \sum_{j \leq i} a_j q_j(r_1, \dots, r_j)] \\ &= b - \sum_{j \leq i} a_j \mathbb{E}_{r_j}[0.5W_j^*(r_j)] = 0.5b \geq 0.5W_{i+1}^*(r_{i+1}) \quad \forall r_{i+1}. \end{aligned}$$

Hence, via induction and Lemma 8, we conclude that $0.5W^*$ is implementable. \square

4.3 Results

4.3.1 Polymatroids

This section considers the gambler who has to choose his strategy from a polymatroid. The gambler's feasible set of strategies can be described as follows:

$$\sum_{j \in I} q_j(r_1, \dots, r_j) \leq g(I) \quad \forall r_1, \dots, r_n, \quad \forall I \subseteq [n],$$

where $[n]$ denotes the set $\{1, \dots, n\}$, and $g : 2^{[n]} \mapsto \mathbb{R}$ is a nondecreasing submodular function.

We now introduce some notation from the literature on generalized permutohedrons,

which contains the polymatroids. A generalized permutohedron is a polytope defined as follows:

$$P_n(\{z_I\}) = \{(q_1, \dots, q_n) \in \mathbb{R}^n : \sum_{j=1}^n q_j = z_{[n]}, \sum_{j \in I} q_j \geq z_I \forall I \subseteq [n]\},$$

where $z_\emptyset = 0$, and z_I is a real number for each $I \subseteq [n]$. Given a polymatroid P associated with g , we can define a generalized permutohedron that includes P . To this end, we first raise the dimension of P to \mathbb{R}^{n+1} through another polytope:

$$\bar{P} := \{(q_1, \dots, q_{n+1}) \in \mathbb{R}^{n+1} : q_{n+1} = g([n]) - \sum_{j=1}^n q_j, (q_1, \dots, q_n) \in P\}.$$

We also define \bar{g} as $\bar{g}(I) := g(I)$ if $I \subseteq [n]$, and $\bar{g}(I) := g([n])$ if $I \ni n+1$. It is easy to check that \bar{g} is a nondecreasing submodular function, and \bar{P} is a generalized permutohedron:

$$\bar{P} = P_{n+1}(\{z_I\}),$$

where $z_{[n]} = \bar{g}([n])$, and $z_I = \bar{g}([n]) - \bar{g}([n] \setminus I)$ for each $I \subseteq [n]$. Note that raising the dimension of P does not necessarily change the optimal solution to the stopping problem, as we can set all rewards of the box $n+1$ to zero.

We need to introduce two more concepts before stating our result. The Minkowski sum of polytopes P and Q in \mathbb{R}^n is defined as $P + Q = \{p + q : p \in P, q \in Q\}$. Moreover, the Minkowski difference $P - Q$ is equal to R if we have $Q + R = P$. Finally, we let Δ denote the standard unit $(n-1)$ -simplex:

$$\Delta = \{(q_1, \dots, q_n) \in \mathbb{R}_+^n : \sum_{j=1}^n q_j = 1\} = \text{conv}\{e_1, \dots, e_n\},$$

where e_i is a binary vector of size n with the only one appearing in the i^{th} column. Also

let the faces of Δ be denoted by

$$\Delta_I = \text{conv}\{e_j : j \in I\} \quad \forall I \subseteq [n].$$

The following result from the generalized permutohedron literature will prove to be helpful.

Proposition 12 (Ardila et al. [44]). *Every generalized permutohedron $P_n(\{z_I\})$ can be written uniquely as a signed Minkowski sum of simplices:*

$$P_n(\{z_I\}) = \sum_{I \subseteq [n]} y_I \Delta_I,$$

where $y_I = \sum_{J \subseteq I} (-1)^{|I|-|J|} z_J$ for each $I \subseteq [n]$.

We utilize Proposition 11 and Proposition 12 to obtain a $\frac{1}{2}$ -prophet inequality for a particular class of generalized permutohedron.

Theorem 9. *Consider a stopping problem with independently drawn nonnegative rewards where the feasible online policies live in a generalized permutohedron $P_n(\{z_I\}) \subseteq \mathbb{R}^n$. If $P_n(\{z_I\})$ has only nonnegative coefficients in its unique signed Minkowski sum representation, the optimal online strategy achieves at least $\frac{1}{2}$ of the prophet's value.*

Proof. Let $w^* \in P_n(\{z_I\})$ be the optimal offline solution. As any $P_n(\{z_I\})$ has a unique signed Minkowski sum representation due to Proposition 12, there must exist $w_I^* \in \Delta_I$ for all $I \subseteq [n]$ such that $w^* = \sum_{I \subseteq [n]} y_I w_I^*$.

Assume that the signed Minkowski sum of $P_n(\{z_I\})$ only has non-negative coefficients, i.e., $y_I \geq 0$ for all $I \subseteq [n]$. Then, as the objective function of the problem is linear over w^* , we can choose each w_I^* to be the optimal offline solution of the problem when the prophet is limited to Δ_I . If we let $Z_{\text{off}}(P)$ denote the optimal objective value of the

offline problem under some polyhedron P , then the following equality must hold:

$$Z_{off}(P_n(\{z_I\})) = \sum_{I \subseteq [n]} y_I Z_{off}(\Delta_I).$$

Notice that any simplex Δ and all its faces contain only one constraint. Therefore, utilizing Proposition 11, we deduce that they should admit $\frac{1}{2}$ -prophet inequality. Letting $Z_{on}(P)$ denote the optimal objective value of the online problem under some polyhedron P , this translates as:

$$0.5Z_{off}(\Delta_I) \leq Z_{on}(\Delta_I) \quad \forall I \subseteq [n].$$

Therefore, we can conclude our proof using the above prophet inequalities:

$$0.5Z_{off}(P_n(\{z_I\})) = 0.5 \sum_{I \subseteq [n]} y_I Z_{off}(\Delta_I) \leq \sum_{I \subseteq [n]} y_I Z_{on}(\Delta_I) \leq Z_{on}(P_n(\{z_I\})),$$

where the last inequality follows from the fact that the Minkowski sum of the optimal online solutions $q_I^* \in \Delta_I$, $I \subseteq [n]$ is again an online solution, and it is an element of $P_n(\{z_I\})$. \square

Theorem 9 provides a simple way to derive $\frac{1}{2}$ -prophet inequalities for a class of polymatroids, as well as the intuition behind this result. Moreover, it gives an easy roadmap to construct an online strategy with half the optimal offline value. When the polymatroid has a Minkowski sum representation with only nonnegative coefficients, we can focus on simplices, which admit easy-to-use online strategies from literature with half the optimal offline value, and take their Minkowski sum.

To underline the potential of our approach, we summarize the results of Kleinberg & Weinberg [41], who consider a stopping problem under a matroid constraint. Kleinberg & Weinberg [41] first introduce a property for some $\alpha > 1$ and say that any deterministic monotone algorithm satisfying this property has α -balanced thresholds. The property ensures that the gambler does not lose much when he sticks to the algorithm. Then, their first

result states that following a monotone algorithm with α -balanced thresholds, the gambler gains at least $\frac{1}{\alpha}$ of the prophet's gain. To prove their main result, *i.e.*, the prophet inequality for $\alpha = 2$, they introduce an algorithm and show that it has 2-balanced thresholds. Specifically, they use some basic properties of matroids to show that their algorithm satisfies the property for balanced thresholds when $\alpha = 2$.

4.3.2 Correlated Rewards

Notation: In this subsection and the following, we denote the set of rewards as \mathcal{R} . Also, \vec{r} denotes the vector of rewards (r_1, \dots, r_i) , to which we also refer as (r_j, \vec{r}_{-j}) given any state $j = 1, \dots, i$. Here, \vec{r}_{-j} is obtained from removing the j^{th} reward from \vec{r} .

In this section, we drop the assumption of independent rewards and reproduce the $\frac{1}{n}$ -prophet inequality result from Hill & Kertz [37]. To this end, we will prove feasibility of the following system of inequalities for all $i + 1 \leq n$.

$$\begin{aligned}
\mathbb{E}_{\vec{r}}[z(\vec{r})] &\geq \frac{1}{i+1} \min_k b_k \\
z(\vec{r}) + \sum_{j \leq i} a_{kj} w_j(\vec{r}) &\leq b_k && \forall \vec{r}, \forall k \leq K, \\
w_j(\vec{r}) &= q_j(r_1, \dots, r_j) && \forall \vec{r}, \forall j \leq i, \\
\mathbb{E}_{\vec{r}_{-j} | r_j} [w_j(\vec{r})] &\geq \frac{1}{i+1} W_j(r_j) && \forall r_j, \forall j \leq i, \\
z(\vec{r}) &\geq 0 && \forall \vec{r}, \\
w_j(\vec{r}) &\geq 0 && \forall \vec{r}, \forall j \leq i, \\
q_j(r_1, \dots, r_j) &\geq 0 && \forall r_1, \dots, r_j, \forall j \leq i.
\end{aligned}$$

Without loss of generality, we can assume that z variables have coefficient 1 in each constraint, as we can either ignore the constraint (when the coefficient is zero) or scale it by the coefficient of z . We work with w_j variables in this model to avoid any confusion due to

dependence between states. The third constraint ensures that for any $j \leq i$, $w_j(r_1, \dots, r_i)$ variables have the same value for all r_{j+1}, \dots, r_i . In other words, w agrees with some online rule q . These constraints are called non-anticipativity constraints in the literature. Furthermore, due to dependence, $W_j(r_j)$ is now defined as $\sum_{\vec{r}_{-j}} f(\vec{r}_{-j} | r_j) w_j(\vec{r})$. The above model can also be written as follows:

$$\begin{aligned}
\sum_{\vec{r}} f(\vec{r}) z(\vec{r}) &\geq \frac{1}{i+1} \min_k b_k \\
f(\vec{r}) \left[z(\vec{r}) + \sum_{j \leq i} a_{kj} w_j(\vec{r}) \right] &\leq f(\vec{r}) b_k && \forall \vec{r}, \forall k \leq K, \\
f(\vec{r}) w_j(\vec{r}) &= f(\vec{r}) q_j(r_1, \dots, r_j) && \forall \vec{r}, \forall j \leq i, \\
\sum_{\vec{r}_{-j}} \frac{f(r_j, \vec{r}_{-j})}{f(r_j)} w_j(r_j, \vec{r}_{-j}) &\geq \frac{1}{i+1} W_j(r_j) && \forall r_j, \forall j \leq i, \\
z(\vec{r}) &\geq 0 && \forall \vec{r}, \\
w_j(\vec{r}) &\geq 0 && \forall \vec{r}, \forall j \leq i, \\
q_j(r_1, \dots, r_j) &\geq 0 && \forall r_1, \dots, r_j, \forall j \leq i.
\end{aligned} \tag{4.7}$$

Theorem 10. Consider a stopping problem with nonnegative rewards where the feasible online policies live in a polymatroid described by $A \in \mathbb{R}_+^{K \times n}$ and $b \in \mathbb{R}^n$. The optimal online strategy achieves at least $\frac{1}{n}$ of the prophet's value.

Proof. If the model in (4.7) is feasible, then its Farkas' alternative must be infeasible:

$$f(\vec{r}) \left[\sum_k \mu_k(\vec{r}) - \tau \right] \geq 0 \quad \forall \vec{r}, \tag{4.8}$$

$$\sum_k a_{kj} f(\vec{r}) \mu_k(\vec{r}) + f(\vec{r}) \beta_j(\vec{r}) - \frac{f(r_j, \vec{r}_{-j})}{f(r_j)} \lambda_j(r_j) \geq 0 \quad \forall \vec{r}, \forall j \leq i, \tag{4.9}$$

$$- \sum_{r_{j+1}, \dots, r_i} f(r_1, \dots, r_i) \beta_j(r_1, \dots, r_i) \geq 0 \quad \forall r_1, \dots, r_j, \forall j \leq i, \tag{4.10}$$

$$\sum_k \sum_{\vec{r}} b_k f(\vec{r}) \mu_k(\vec{r}) - \sum_{j \leq i} \sum_{r_j} \frac{W_j(r_j)}{i+1} \lambda_j(r_j) - \tau \frac{\min_k b_k}{i+1} < 0, \tag{4.11}$$

where all variables, except β , are nonnegative. We want to show that the above alternative is infeasible. That is, all nonnegative (μ, λ, τ) that satisfy the first three constraints lead to:

$$\sum_k b_k \mathbb{E}_{\vec{r}}[\mu_k(\vec{r})] - \sum_{j \leq i} \sum_{r_j} \frac{W_j(r_j)}{i+1} \lambda_j(r_j) - \tau \frac{\min_k b_k}{i+1} \geq 0,$$

contradicting (4.11).

First, notice that β variables are free and do not appear in (4.11). Moreover, decreasing β variables do not violate (4.10) so that we can focus on the solutions where (4.9) are always binding:

$$f(\vec{r}) \beta_j(\vec{r}) = \frac{f(r_j, \vec{r}_{-j})}{f(r_j)} \lambda_j(r_j) - \sum_k a_{kj} f(\vec{r}) \mu_k(\vec{r}) \quad \forall \vec{r}, \forall j \leq i.$$

Hence, the constraint (4.10) becomes:

$$\sum_{r_{j+1}, \dots, r_i} \sum_k a_{kj} f(\vec{r}) \mu_k(\vec{r}) - \sum_{r_{j+1}, \dots, r_i} \frac{f(r_j, \vec{r}_{-j})}{f(r_j)} \lambda_j(r_j) \geq 0 \quad \forall r_1, \dots, r_j, \forall j \leq i. \quad (4.10')$$

Using this new constraint, we can upper bound λ variables as follows. We first sum (4.10') over r_1, \dots, r_{j-1} :

$$\begin{aligned} \sum_k \sum_{r_{j+1}, \dots, r_i} a_{kj} f(\vec{r}) \mu_k(\vec{r}) &\geq \sum_{r_{j+1}, \dots, r_i} \frac{f(r_j, \vec{r}_{-j})}{f(r_j)} \lambda_j(r_j) \quad \forall r_1, \dots, r_j, \forall j \leq i, \\ \sum_k \sum_{\vec{r}_{-j}} a_{kj} f(r_j, \vec{r}_{-j}) \mu_k(r_j, \vec{r}_{-j}) &\geq \sum_{\vec{r}_{-j}} \frac{f(r_j, \vec{r}_{-j})}{f(r_j)} \lambda_j(r_j) \quad \forall r_j, \forall j \leq i. \end{aligned}$$

As $\sum_{\vec{r}_{-j}} f(r_j, \vec{r}_{-j}) = f(r_j)$, we can rewrite the above upper bound as follows:

$$\begin{aligned} \sum_k f(r_j) \sum_{\vec{r}_{-j}} a_{kj} \frac{f(r_j, \vec{r}_{-j})}{f(r_j)} \mu_k(r_j, \vec{r}_{-j}) &\geq \lambda_j(r_j) && \forall r_j, \forall j \leq i, \\ \sum_k f(r_j) a_{kj} \mathbb{E}_{\vec{r}_{-j}}[\mu_k(r_j, \vec{r}_{-j}) \mid r_j] &\geq \lambda_j(r_j) && \forall r_j, \forall j \leq i. \end{aligned}$$

Substituting λ variables with these upper bounds gives us the following lower bound for the left-hand side of (4.11):

$$\sum_k b_k \mathbb{E}_{\vec{r}}[\mu_k(\vec{r})] - \sum_{j \leq i} \sum_{r_j} \frac{W_j(r_j)}{i+1} \sum_k f(r_j) a_{kj} \mathbb{E}_{\vec{r}_{-j}}[\mu_k(r_j, \vec{r}_{-j}) \mid r_j] - \tau \frac{\min_k b_k}{i+1}. \quad (4.11')$$

Now, we go back to the offline problem to develop a lower bound for the second term of (4.11'). Ignore the rewards after period i and assume that we only solve for the first i random variables. Let \hat{w} denote the optimal solution of this problem and let $\hat{W}_j(r_j)$ denote its expectation over \vec{r}_{-j} given reward r_j . We multiply both sides of all constraints by $f(\vec{r})$ and sum them over \vec{r}_{-j} :

$$\begin{aligned} a_{kj} \hat{w}_j(\vec{r}) &\leq b_k && \forall \vec{r}, \forall j \leq i, \forall k \leq K, \\ f(r_j) \frac{f(\vec{r})}{f(r_j)} a_{kj} \hat{w}_j(\vec{r}) &\leq f(\vec{r}) b_k && \forall \vec{r}, \forall j \leq i, \forall k \leq K, \\ f(r_j) a_{kj} \hat{W}_j(r_j) &\leq f(r_j) b_k && \forall r_j, \forall j \leq i, \forall k \leq K. \end{aligned}$$

Notice that since we ignore the rewards after period i , for any $j \leq i$, optimal values of \hat{w} should always be bigger than the optimal solution of the original offline problem. Hence, we have $W_j(r_j) \leq \hat{W}_j(r_j)$ for all r_j and $j \leq i$. Then, we can substitute $\hat{W}_j(r_j)$ from the left-hand side of the above inequality with $W_j(r_j)$. To obtain the promised lower bound, we multiply both sides of the above inequality by $\mathbb{E}_{\vec{r}_{-j}}[\mu_k(r_j, \vec{r}_{-j}) \mid r_j]$, which are all

nonnegative. As a result, we obtain the following inequalities for all $r_i, j \leq i$, and $k \leq K$:

$$f(r_j)a_{kj}W_j(r_j)\mathbb{E}_{\vec{r}_{-j}}[\mu_k(r_j, \vec{r}_{-j}) \mid r_j] \leq f(r_j)b_k\mathbb{E}_{\vec{r}_{-j}}[\mu_k(r_j, \vec{r}_{-j}) \mid r_j].$$

Summing these inequalities over r_j yields:

$$\sum_{r_j} f(r_j)a_{kj}W_j(r_j)\mathbb{E}_{\vec{r}_{-j}}[\mu_k(r_j, \vec{r}_{-j}) \mid r_j] \leq b_k\mathbb{E}_{\vec{r}}[\mu_k(\vec{r})] \quad \forall j \leq i, \forall k \leq K.$$

Finally, we sum them over $j \leq i$ and $k \leq K$ to obtain a lower bound for the second term of (4.11').

$$\sum_{j \leq i} \sum_k \sum_{r_j} f(r_j)a_{kj}W_j(r_j)\mathbb{E}_{\vec{r}_{-j}}[\mu_k(r_j, \vec{r}_{-j}) \mid r_j] \leq i \sum_k b_k\mathbb{E}_{\vec{r}}[\mu_k(\vec{r})].$$

Thus, (4.11') is bounded below by:

$$\sum_k b_k\mathbb{E}_{\vec{r}}[\mu_k(\vec{r})] - \frac{i}{i+1} \sum_k b_k\mathbb{E}_{\vec{r}}[\mu_k(\vec{r})] - \tau \frac{\min_k b_k}{i+1}.$$

The following arguments verify that this lower bound is nonnegative:

$$\begin{aligned} \frac{1}{i+1} \sum_k b_k\mathbb{E}_{\vec{r}}[\mu_k(\vec{r})] - \tau \frac{\min_k b_k}{i+1} &\geq \frac{1}{i+1} \min_k b_k \sum_{k'} \mathbb{E}_{\vec{r}}[\mu_{k'}(\vec{r})] - \tau \frac{\min_k b_k}{i+1}, \\ &\geq \frac{1}{i+1} \min_k b_k \left[\sum_{k'} \mathbb{E}_{\vec{r}}[\mu_{k'}(\vec{r})] - \tau \right], \\ &= \frac{1}{i+1} \min_k b_k \left[\mathbb{E}_{\vec{r}} \left[\sum_{k'} \mu_{k'}(\vec{r}) \right] - \tau \right] \geq 0, \end{aligned}$$

where the last inequality follows due to (4.8). This observation lets us conclude that the considered Farkas' alternative is infeasible, which means:

$$h_{i+1} \left(\frac{W_1}{i+1}, \dots, \frac{W_i}{i+1} \right) = \max \mathbb{E}_{\vec{r}}[z(\vec{r})] \geq \frac{1}{i+1} \min_k b_k \geq \frac{W_{i+1}(r_{i+1})}{i+1} \quad \forall r_{i+1}.$$

Hence, via induction and Lemma 8, we conclude that $\frac{W^*}{n}$ is implementable. \square

4.3.3 K Many Constraints

In this section, we prove that any stopping problem with independently drawn rewards and K many constraints admits $\frac{1}{K+1}$ -prophet inequality. As in Subsection 4.3.2, we will show that $h_{i+1}(\frac{1}{K+1}W_1^*, \dots, \frac{1}{K+1}W_i^*) \geq \frac{1}{K+1}W_{i+1}^*(r_{i+1})$ for all r_{i+1} .

Theorem 11. *Consider a stopping problem with independently drawn nonnegative rewards where the feasible online policies live in a polymatroid described by $A \in \mathbb{R}_+^{K \times n}$ and $b \in \mathbb{R}^n$. The optimal online strategy achieves at least $\frac{1}{K+1}$ of the prophet's value.*

Proof. First, notice that $(\frac{1}{K+1}W_1^*, \dots, \frac{1}{K+1}W_{K+1}^*)$ is feasible thanks to scaling by $\frac{1}{K+1}$. Hence, as an induction hypothesis, we can assume that there exists some $i < n$ such that:

$$h_j(\frac{1}{K+1}W_1^*, \dots, \frac{1}{K+1}W_{j-1}^*) \geq \frac{1}{K+1}W_j^*(r_j) \quad \forall r_j, \forall j \leq i.$$

Then, the problem $h_{i+1}(\frac{1}{K+1}W_1^*, \dots, \frac{1}{K+1}W_i^*)$ is feasible, and for any feasible \vec{q} with the interim values $(\frac{1}{K+1}W_1^*, \dots, \frac{1}{K+1}W_i^*)$, it is optimal to set:

$$z^*(\vec{r}) = \min_k \{b_k - \sum_{j \leq i} a_{kj} q_j(r_1, \dots, r_j)\} \quad \forall \vec{r}.$$

We partition the set of rewards \mathcal{R} into K sets with respect to the tightest bound on $z^*(\vec{r})$ variables. Define a family of subsets $(\mathcal{R}_k)_{k=1}^K$ as follows:

$$\mathcal{R}_k = \{\vec{r} \in \mathcal{R} \mid z^*(\vec{r}) = b_k - \sum_{j \leq i} a_{kj} q_j(r_1, \dots, r_j)\} \quad \forall k = 1, \dots, K,$$

such that $\cup_{k=1}^K \mathcal{R}_k = \mathcal{R}$, and $\mathcal{R}_k \cap \mathcal{R}_l = \emptyset$ for all $k, l \in \{1, \dots, K\}$ with $k \neq l$.

We can use $(\mathcal{R}_k)_{k=1}^K$ to write the objective function as follows:

$$h_{i+1}(0.5W_1^*, \dots, 0.5W_i^*) = \mathbb{E}_{\vec{r}}[z^*(\vec{r})] = \sum_{k=1}^K \sum_{\vec{r} \in \mathcal{R}_k} f(\vec{r}) [b_k - \sum_{j \leq i} a_{kj} q_j(r_1, \dots, r_j)].$$

Finally, we define probability $f_k = \sum_{\vec{r} \in \mathcal{R}_k} f(\vec{r})$ for all $k = 1, \dots, K$ and set $\mathcal{K} = \{k \mid f_k \geq \frac{1}{K+1}\}$. Notice that \mathcal{K} cannot be empty since otherwise, we would have the following contradiction:

$$\sum_{k=1}^K f_k < \sum_{k=1}^K \frac{1}{K+1} < 1.$$

Then, the objective has the following lower bound:

$$\begin{aligned} h_{i+1}(0.5W_1^*, \dots, 0.5W_i^*) &\geq \sum_{k \in \mathcal{K}} f_k b_k - \sum_{k \in \mathcal{K}} \sum_{\vec{r} \in \mathcal{R}_k} f(\vec{r}) \sum_{j \leq i} a_{kj} q_j(r_1, \dots, r_j), \\ &\geq \sum_{k \in \mathcal{K}} f_k b_k - \sum_{k \in \mathcal{K}} \sum_{\vec{r} \in \mathcal{R}} f(\vec{r}) \sum_{j \leq i} a_{kj} q_j(r_1, \dots, r_j), \\ &= \sum_{k \in \mathcal{K}} f_k b_k - \sum_{k \in \mathcal{K}} \sum_{j \leq i} a_{kj} \sum_{r_j} f(r_j) \frac{W_j(r_j)}{K+1}. \end{aligned}$$

We know from the offline problem that $\sum_{j \leq i} a_{kj} \sum_{r_j} f(r_j) \frac{W_j(r_j)}{K+1} \leq \frac{b_k}{K+1}$. Hence, the following lower bound follows:

$$\begin{aligned} h_{i+1}(0.5W_1^*, \dots, 0.5W_i^*) &\geq \sum_{k \in \mathcal{K}} \left(f_k - \frac{1}{K+1} \right) b_k \geq \sum_{k \in \mathcal{K}} \left(f_k - \frac{1}{K+1} \right) \min_{l \in \{1, \dots, K\}} b_l \\ &= \left(\sum_{k \in \mathcal{K}} f_k - \frac{|\mathcal{K}|}{K+1} \right) \min_{l \in \{1, \dots, K\}} b_l \\ &= \left(\frac{K+1 - |\mathcal{K}|}{K+1} - \sum_{k \notin \mathcal{K}} f_k \right) \min_{l \in \{1, \dots, K\}} b_l > \frac{1}{K+1} \min_{l \in \{1, \dots, K\}} b_l \\ &\geq \frac{1}{K+1} W_{i+1}^*(r_{i+1}) \quad \forall r_{i+1}, \end{aligned}$$

where the penultimate inequality follows from the fact that $\sum_{k \notin \mathcal{K}} f_k < \frac{K-|\mathcal{K}|}{K+1}$. Hence, via

induction and Lemma 8, we conclude that $\frac{W^*}{K+1}$ is implementable. □

4.4 Concluding Remarks

By utilizing the extensive tools in the linear programming literature, we more straightforwardly reproduced various results from the prophet inequality literature. Hence, our methodology proves to be fruitful and promising. Here we note that our result in Subsection 4.3.1 is valid only for a subset of polymatroids. Specifically, we show that a polymatroid admits $\frac{1}{2}$ -prophet inequality if its unique signed Minkowski sum only has nonnegative coefficients. As the $\frac{1}{2}$ -prophet inequality result is already proven for all polymatroids by Dütting & Kleinberg [42], we conjecture that our proof extends to all polymatroids. Hopefully, our results will provide new insights and eventually lead to novel prophet inequalities and robust mechanism designs.

Bibliography

- [1] A. Attar, T. Mariotti, and F. Salanié, “Nonexclusive competition in the market for lemons,” *Econometrica*, vol. 79, no. 6, pp. 1869–1918, 2011.
- [2] A. Attar, T. Mariotti, and F. Salanié, “Nonexclusive competition under adverse selection,” *Theoretical Economics*, vol. 9, no. 1, pp. 1–40, 2014.
- [3] E. Ben-Porath, E. Dekel, and B. L. Lipman, “Optimal allocation with costly verification,” *American Economic Review*, vol. 104, no. 12, pp. 3779–3813, 2014.
- [4] Y. Li, “Mechanism design with costly verification and limited punishments,” *Journal of Economic Theory*, vol. 186, p. 105000, 2020.
- [5] T. Mylovanov and A. Zapechelnyuk, “Optimal allocation with ex post verification and limited penalties,” *American Economic Review*, vol. 107, no. 9, pp. 2666–94, 2017.
- [6] E. Delage and Y. Ye, “Distributionally robust optimization under moment uncertainty with application to data-driven problems,” *Operations Research*, vol. 58, no. 3, pp. 595–612, 2010.
- [7] W. Wiesemann, D. Kuhn, and M. Sim, “Distributionally robust convex optimization,” *Operations Research*, vol. 62, no. 6, pp. 1358–1376, 2014.
- [8] D. A. Iancu and N. Trichakis, “Pareto efficiency in robust optimization,” *Management Science*, vol. 60, no. 1, pp. 130–147, 2014.

- [9] R. M. Townsend, “Optimal contracts and competitive markets with costly state verification,” *Journal of Economic Theory*, vol. 21, no. 2, pp. 265–293, 1979.
- [10] G. A. Chua, G. Hu, and F. Liu, “Optimal multi-unit allocation with costly verification,” *Available at SSRN 3407031*, 2019.
- [11] H. I. Bayrak, K. Güler, and M. Ç. Pınar, “Optimal allocation with costly inspection and discrete types under ambiguity,” *Optimization Methods and Software*, vol. 32, no. 4, pp. 699–718, 2017.
- [12] S. R. Balseiro, A. Kim, and D. Russo, “On the futility of dynamics in robust mechanism design,” *Operations Research*, vol. 69, no. 6, pp. 1767–1783, 2021.
- [13] V. Carrasco, V. F. Luz, N. Kos, M. Messner, P. Monteiro, and H. Moreira, “Optimal selling mechanisms under moment conditions,” *Journal of Economic Theory*, vol. 177, pp. 245–279, 2018.
- [14] Z. Chen, Z. Hu, and R. Wang, “Screening with limited information: The minimax theorem and a geometric approach,” *Available at SSRN 3940212*, 2021.
- [15] Ç. Koçyiğit, N. Rujeeapaiboon, and D. Kuhn, “Robust multidimensional pricing: Separation without regret,” *Mathematical Programming*, pp. 1–34, 2021.
- [16] M. Ç. Pınar and C. Kızılkale, “Robust screening under ambiguity,” *Mathematical Programming*, vol. 163, no. 1, pp. 273–299, 2017.
- [17] S. Wang, S. Liu, and J. Zhang, “Minimax regret mechanism design with moments information,” *Available at SSRN 3707021*, 2020.
- [18] J. Anunrojwong, S. Balseiro, and O. Besbes, “On the robustness of second-price auctions in prior-independent mechanism design,” *arXiv preprint arXiv:2204.10478*, 2022.
- [19] W. He and J. Li, “Correlation-robust auction design,” *Journal of Economic Theory*, vol. 200, p. 105403, 2022.

- [20] Ç. Koçyiğit, G. Iyengar, D. Kuhn, and W. Wiesemann, “Distributionally robust mechanism design,” *Management Science*, vol. 66, no. 1, pp. 159–189, 2020.
- [21] A. Suzdaltsev, “An optimal distributionally robust auction,” *arXiv preprint arXiv:2006.05192*, 2020.
- [22] D. Walton and G. Carroll, “A general framework for robust contracting models,” *Econometrica*, 2022.
- [23] D. Fudenberg and J. Tirole, *Game Theory*. MIT Press, 1991.
- [24] G. Akerlof, “The market for “lemons”: Quality uncertainty and the market mechanism,” *Quarterly Journal of Economics*, vol. 84, no. 3, pp. 488–500, 1970.
- [25] M. Rothschild and J. Stiglitz, “Equilibrium in competitive insurance markets: An essay on the economics of imperfect information,” *Quarterly Journal of Economics*, vol. 90, no. 4, pp. 629–649, 1976.
- [26] A. Mas-Colell, M. D. Whinston, and J. R. Green, *Microeconomic Theory*. New York: Oxford University Press, 1995.
- [27] R. Inderst and A. Wambach, “Competitive insurance markets under adverse selection and capacity constraints,” *European Economic Review*, vol. 45, no. 10, pp. 1981–1992, 2001.
- [28] R. Inderst and A. Wambach, “Capacity constrained firms in (labor) markets with adverse selection,” *Economic Theory*, vol. 19, no. 3, pp. 525–548, 2002.
- [29] M. Peters, “Common agency and the revelation principle,” *Econometrica*, vol. 69, no. 5, pp. 1349–1372, 2001.
- [30] D. Martimort and L. Stole, “The revelation and delegation principles in common agency games,” *Econometrica*, vol. 70, no. 4, pp. 1659–1673, 2002.

- [31] A. Attar, T. Mariotti, and F. Salanié, “On competitive nonlinear pricing,” *Theoretical Economics*, vol. 14, no. 1, pp. 297–343, 2019.
- [32] A. Pelletier and C. Thomas, “Information in online labour markets,” *Oxford Review of Economic Policy*, vol. 34, no. 3, pp. 376–392, 2018.
- [33] A. Pallais, “Inefficient hiring in entry-level labor markets,” *American Economic Review*, vol. 104, no. 11, pp. 3565–3599, 2014.
- [34] C. T. Stanton and C. Thomas, “Landing the first job: The value of intermediaries in online hiring,” *Review of Economic Studies*, vol. 83, no. 2, pp. 810–854, 2016.
- [35] U. Krengel and L. Sucheston, “Semiamarts and finite values,” *Bulletin of the American Mathematical Society*, vol. 83, no. 4, pp. 745–747, 1977.
- [36] M. Epitropou and R. Vohra, “Optimal on-line allocation rules with verification,” in *International Symposium on Algorithmic Game Theory*, pp. 3–17, Springer, 2019.
- [37] T. P. Hill and R. P. Kertz, “A survey of prophet inequalities in optimal stopping theory,” *Contemp. Math*, vol. 125, pp. 191–207, 1992.
- [38] M. T. Hajiaghayi, R. Kleinberg, and T. Sandholm, “Automated online mechanism design and prophet inequalities,” in *AAAI*, vol. 7, pp. 58–65, 2007.
- [39] S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan, “Multi-parameter mechanism design and sequential posted pricing,” in *Proceedings of the forty-second ACM symposium on Theory of computing*, pp. 311–320, 2010.
- [40] S. Alaei, “Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers,” *SIAM Journal on Computing*, vol. 43, no. 2, pp. 930–972, 2014.
- [41] R. Kleinberg and S. M. Weinberg, “Matroid prophet inequalities,” in *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pp. 123–136, 2012.

- [42] P. Dütting and R. Kleinberg, “Polymatroid prophet inequalities,” in *Algorithms-ESA 2015*, pp. 437–449, Springer, 2015.
- [43] B. Lucier, “An economic view of prophet inequalities,” *ACM SIGecom Exchanges*, vol. 16, no. 1, pp. 24–47, 2017.
- [44] F. Ardila, C. Benedetti, and J. Doker, “Matroid polytopes and their volumes,” *Discrete & Computational Geometry*, vol. 43, no. 4, pp. 841–854, 2010.