Electronic Service Allocation with Private Quality Information

Dipl.-Math. Tobias Widmer

DISSERTATION

Submitted to the Faculty of Business, Economics, and Social Sciences at
the University of Hohenheim

In Partial Fulfillment of the Requirements for the Degree of
Doctor of Economics (Dr. oec.)

June 2017

Date of the Oral Examination:	June 28, 2017

Principal Advisor:	Prof. Dr. Stefan Kirn
Co-Advisor:	Prof. Dr. Rainer Alt
Chair of the Examination Committee:	Prof. Dr. Jörg Schiller
Dean of the Faculty:	Prof. Dr. Dirk Hachmeister

Abstract

The efficient allocation of electronic services is a complex business problem. Customers demand electronic services from service providers who supply these services at a specified quality of service (QoS). Electronic marketplaces provide a platform on which multiple customers and multiple providers negotiate the allocation of electronic services. Such marketplaces might be administrated by government authorities or large corporations who aim at a socially optimal allocation. If, however, the QoS desired by customers and the QoS offered by providers is private information on both market sides, it is difficult to design mechanisms that result in an optimal allocation from the perspective of a social planner. Using a mechanism design framework, this research studies the allocation of electronic services with private quality information. Because private information is present in the analysis, a second-best allocation mechanism is derived that satisfies incentive compatibility, individual rationality, and budget balance. The objectives of this research are (i) to develop a double-sided mechanism for allocating electronic services with private quality information and (ii) study this mechanism's efficiency properties in a set of simulation experiments to demonstrate its usefulness. All experiments imply that the asymptotic efficiency of the mechanism is bounded away from 100% even for large markets. This finding is related to the economic concept of informational smallness that arises in this framework. Because demand and supply are characterized by services of distinct quality, none of the traders become informationally small as the market size increases.

Contents Overview

Contents

List of Figures

List of Tables

List of Symbols

a_i	Customer i
$\mathscr{A}_C$	Set of customers
b_j	Provider j
$\mathscr{A}_P$	Set of providers
$c(\theta_i, \sigma_j)$	Provision cost of b_j
$\mathbb{E}$	Unconditional expectation
$\mathbb{E}_{\theta_i}$	Expectation conditional on type θ_i
$f_{sc} : \Theta_1 \times \cdots \times \Theta_I \to \Omega$	Social choice function
f_i	Probability density function of customer a_i's type θ_i
$f_{(k)}$	Probability density of the k-th order statistic for customers
F_i	Probability distribution of customer a_i's type θ_i
$F_{(k)}$	Cumulative distribution associated with $f_{(k)}$
g	Outcome function of a mechanism
h_j	Probability density function of provider b_j's type σ_j
$h_{(k)}$	Probability density of the k-th order statistic for providers
H_j	Probability distribution of provider b_j's type σ_j
$H_{(k)}$	Cumulative distribution associated with $h_{(k)}$
K	Minimum of N and M
M	Number of providers
$\mathscr{M}$	Mechanism
N	Number of customers

Ω	Set of all possible allocation alternatives
$\psi_C(\theta_i, \sigma_j)$	Virtual valuation of customer a_i
$\psi_P(\theta_i, \sigma_j)$	Virtual provision cost of provider b_j
$\rho_\sigma(\sigma_j)$	Rank of σ_j within all provider types
$\rho_\theta(\theta_i)$	Rank of θ_i within all customer types
$R_C(\theta_i)$	Reserve function of customer a_i
$R_P(\sigma_j)$	Reserve function of provider b_j
σ	Type vector of all providers
σ_j	Type of provider b_j
σ_{-j}	Type vector of other providers except b_j's type
$\bar{\sigma}_j$	Highest possible type among all providers
$s_i : \Omega_i \to S_I$	Strategy function of a_i
S_i	Strategies of a_i
$S_C^\infty(\theta)$	Informational rent of customer a_i for $K \to \infty$
$S_P^\infty(\sigma)$	Informational rent of provider b_j for $K \to \infty$
θ	Type vector of all customers
θ_i	Type of customer a_i
θ_{-i}	Type vector of other customers except a_i's type
$\underline{\theta}_i$	Lowest possible type among all customers
Θ_i	Type space of a_i
$t_C(\theta_i, \theta_{-i}, \sigma)$	Monetary transfer made by customer a_i, conditional on all other customers' types
$T_C(\theta_i)$	Expected monetary transfer made by customer a_i
$t_P(\theta, \sigma_j, \sigma_{-j})$	Monetary transfer received by provider b_j, conditional on all other providers' types
$T_P(\sigma_j)$	Expected monetary transfer received by provider b_j
$u_i : \Theta_i \times \Omega \to \mathbb{R}$	Utility of a_i

$U_C(\theta_i, \theta_{-i}, \sigma)$	Expected utility of customer a_i, conditional on all other customers' types
$U_P(\theta, \sigma_j, \sigma_{-j})$	Expected utility of provider b_i, conditional on all other providers' types
$v(\theta_i, \sigma_j)$	Valuation of a_i
$w_k^C(\theta_i)$	Probability that θ_i has rank k within all customer types
$w_k^P(\sigma_j)$	Probability that σ_j has rank k within all provider types
x	Allocation alternative
$x_{ij}(\theta, \sigma)$	Probability of allocation between a_i and b_j

Chapter 1

Introduction

"It's no good, it's no good!" says the buyer – then goes off and boasts about the purchase.

Proverbs 20:14, The Bible
(New International Version)

1.1 Motivation

The increasing number of third-party vendors offering electronic services has stimulated the growth of information technology (IT) outsourcing in many organizations (Chang and Gurbaxani, 2012). Firms outsource their applications to these vendors, who provide access to computing resources at a specified quality of service (QoS) such as availability, throughput, and execution time (Gartner, 2013). Advances in IT facilitate the substitution of traditionally static and long-term relationships by flexible contracts of shorter duration, with Cloud computing being the most recent manifestation of these advances (Armbrust et al., 2010). To this end, electronic marketplaces for Cloud services such as the Avnet Cloud Marketplace (2016) have emerged. These marketplaces might be administrated by government authorities or large corporations who aim at a socially optimal allocation. For instance, government users may demand electronic services from geographically remote data centers owned by the government (Kundra, 2011). The overall objective of the marketplace is to serve the customers and to best utilize these data centers.

This research addresses the allocation problem for electronic services with private information about the QoS. On electronic marketplaces, customers and providers automatically negotiate the exchange of these electronic services. By assigning specific valuations to the demands and offers, each trader can enunciate his preferences over the allocation of the services. Once all demands and offers are submitted, the mechanism determines an allocation that maximizes the sum of the aggregated preferences. However, the design of such mechanisms is difficult because of the following requirements:

1. **Double-sided competition:** Electronic marketplaces are platforms that bring together two distinct groups of users to facilitate their interaction (Eisenmann et al., 2006). Surplus is generated when the marketplace matches the demand of customers with the supply of providers. The emerging double-sided competition significantly influences the pricing structure offered to the traders (Armstrong, 2006). Therefore, the allocation mechanism proposed by the platform must maintain double-sided competition between the traders.

2. **QoS awareness:** Customers have different requirements for the quality characteristics of the services (Pal and Hui, 2012). Therefore, it is not viable to solely account for the price of the services as the single attribute. On the other hand, service providers use QoS attributes for differentiation from the competition. Hence, the allocation of electronic service must be predicated on the specific QoS requirements of customers and the distinct QoS offers of providers. However, the offered QoS is typically fixed prior to provider selection (Bichler and Kalagnanam, 2005), and mechanisms that internalize QoS significantly affect the outcome throughout the allocation process (Bockstedt and Goh, 2011). Therefore, the desired QoS of

customers and the actual QoS offered by service providers must be internalized in the allocation mechanism.

3. **Private information:** Customers and providers are not capable of perfectly observing exact information about the demand and supply of QoS on the market. In fact, each customer's desired QoS is known only to that customer, and each provider's actual QoS is known only to that provider. The mechanism designer (e.g., a social planner) observes no one's desired or actual QoS, and no individual observes the QoS of any other individual. Hence, the mechanism must facilitate the allocation of electronic services for which any QoS information is unknown.

4. **Incentive compatibility:** Self-interested customers and providers want to maximize their own expected utility. They may act strategically when negotiating with others if this is advantageous to them. Thus, the mechanism has to provide adequate incentives to strategic individuals because they might potentially misreport their true preferences.

5. **Individual rationality:** Customers and providers may only participate in the allocation mechanism if they can expect non-negative utilities from their participation. If they anticipate negative utility, they will withdraw from the mechanism. Hence, the mechanism must not force individuals to participate in the allocation process.

6. **Budget balance:** The mechanism must omit any independent intermediary in order to facilitate distributed decision-making among the traders (Egri and Váncza, 2013). This requirement implies that all payments must be distributed among the traders.

7. **Optimality:** The ultimate objective of the mechanism is to achieve an outcome that is optimal from a social welfare perspective. In an *ex post* optimal mechanism,

customers receive the electronic services whenever their valuation for the services is higher than the cost of the providers (Myerson and Satterthwaite, 1983).

Standard impossibility theorems from mechanism design theory assert that meeting these requirements simultaneously is not attainable (Laffont and Maskin, 1979; Myerson and Satterthwaite, 1983). In particular, ex post optimality cannot be attained when incentive compatibility, individual rationality, and budget balance are required as well. Therefore, the mechanism designer must determine a viable trade-off of these requirements. One possible compromise in the presence of privately known QoS is to derive a *second-best* mechanism. A second-best mechanism is the optimal mechanism when private information is present. Unlike *first-best* mechanisms, which are optimal ex post, second-best mechanisms only achieve *ex ante* optimality. That is, second-best mechanisms maximize the expected social welfare subject to incentive compatibility, individual rationality, and budget balance. The outcome of such mechanisms can be used to estimate the efficiency loss that must be tolerated in comparison to the first-best outcome (Rustichini et al., 1994).

Studying the efficiency properties of such second-best mechanisms in the presence of private quality information is not possible with current allocation mechanisms. Current approaches for electronic service allocation integrate QoS through multi-attribute procurement auctions with optimal scoring rules (Che, 1993; Bichler and Kalagnanam, 2005; Blau et al., 2010) and combinatorial auctions (Zaman and Grosu, 2013; Karaenke, 2014). However, these models only support competition either among providers or among customers but not both at the same time. The allocation mechanism studied by Regev and Nisan (2000) facilitates double-sided competition but does not explicitly incorporate QoS into the preferences of customers and providers. The mechanism proposed by Schnizler et al. (2008) supports both double-sided competition and QoS integration. Yet their pricing scheme does not satisfy incentive compatibility. In economic theory, Gresik and Satterthwaite (1989) propose an allocation mechanism that supports double-sided competition with private information. Similar to the approach in this thesis, Gresik and Satterthwaite derive a second-best mechanism that satisfies incentive compatibility, individual rationality, and budget balance. However, their model considers the allocation of identical objects only. Insofar, electronic services of distinct QoS cannot be allocated with their approach. Muthoo and Mutuswami (2005, 2011) study the efficiency properties of an allocation mechanism with private information about distinct quality levels of the traded objects. Yet they do not derive the optimal allocation rules or payments used by the second-best mechanism. Johnson (2013) proposes an optimal double-sided matching mechanism, in which the matching of the traders is based on distinct, privately known quality characteristics. Nevertheless, Johnson derives the optimal matching from the perspective of a profit-maximizing intermediary only. Hence, the efficiency properties of the corresponding second-best mechanism cannot be assessed by his approach.

1.2 Research approach

This research studies the allocation of electronic services with private quality information from a mechanism design perspective. Mechanism design theory belongs to the discipline of game theory (Fudenberg and Tirole, 1993). It studies how privately known preferences of multiple players can be aggregated toward a social choice (Nisan and Ronen, 2001). In this research, multiple customers demand electronic services of a specific QoS known to them alone, and multiple providers offer electronic services at a specific QoS known exclusively by them. The prevalent focus of this research is on matching markets, in which potential gains from trade depend on the privately known QoS of the matched customers and providers. Because ex post optimality is unattainable when incentive compatibility, individual rationality, and budget-balance are required, this research derives the optimal allocation rules of a second-best mechanism that maximizes the expected social welfare (i.e., gains from trade) across all customers and providers. Using the first-best outcome as a benchmark, this work studies the efficiency properties of the second-best mechanism subject to incentive compatibility, individual rationality, and budget-balance. To identify these optimal rules, this research focuses on direct revelation mechanisms by invoking the revelation principle (Myerson, 1979) in a first step. Second, the real-world relevance is addressed by implementing the mechanism through a specific *position auction*. For instance, position auctions are used by Google and Yahoo to display search engine ads at positions that are most likely to be clicked on by the user (Varian, 2007). Third, a set of simulation experiments demonstrates the efficacy of the proposed mechanism. Thus, the objectives of this research are to (i) develop a double-sided mechanism for allocating electronic services with privately known QoS and (ii) evaluate this mechanism in a set of simulation experiments to demonstrate its usefulness.

Because the social welfare induced by an allocation depends on each trader's private quality information, this research draws on matching mechanisms, in which customers and providers only produce mutual surplus if they are matched together. Johnson (2013) proposes a double-sided matching mechanism that allocates the traders based on their private quality information. While the approach examined by Johnson (2013) focuses solely on mechanisms that maximize the expected profit of the intermediary, this research studies the efficiency properties of the second-best set of mechanisms for allocating electronic services from a social welfare perspective. Therefore, this research extends the work of Johnson (2013) in two directions. First, instead of maximizing the intermediary's profit, the proposed mechanism maximizes the expected social welfare subject to budget balance. Second, the resulting second-best mechanism is analyzed for its efficiency properties.

1.3 Overview

This thesis is structured in six chapters. The introduction is followed by a literature-based analysis of the state of the art. Then the formal model is presented and applied to the development and evaluation of the allocation mechanism.

Chapter 1 introduces the addressed research problem and provides an overview of the research approach.

Chapter 2 provides the analysis of the state of the art relevant to this research by first discussing the basic concepts of mechanism design theory, followed by inferring a set of requirements on the mechanism. Subsequently, this chapter analyzes the extant literature with respect to the allocation problem.

Chapter 3 introduces the formal model for studying the QoS-aware allocation of electronic services supporting double-sided competition under private quality information.

Chapter 4 presents the allocation mechanism for electronic services with distinct quality values. The mechanism is examined for its efficiency properties from a social welfare perspective.

Chapter 5 reports on the experimental evaluation of the proposed allocation mechanism. A set of simulation experiments with artificial data is conducted to provide evidence of the mechanism's efficacy.

Chapter 6 summarizes the main results and outlines opportunities for future research.

Chapter 2

Analysis of the State of the Art

The simple believe anything, but the prudent give thought to their steps.

Proverbs 14:15, The Bible
(New International Version)

This chapter presents the analysis of the state of the art with respect to the allocation of electronic services with private quality information. Grounded on pertinent fundamentals in mechanism design theory, this chapter derives a set of requirements that must be fulfilled for solving the addressed problem. Subsequently, this chapter analyzes extant allocation mechanisms for electronic services in the literature and assesses to which degree they meet these requirements.

2.1 Mechanism design theory

Mechanism design belongs to the mathematical discipline of game theory, which analyzes how individual agents behave in a given game setting. In this thesis, the terms *agents* and *traders* are used interchangeably. Mechanism design theory is "a strategic version of social choice theory, which adds the assumption that agents will behave so as to maximize their individual payoffs" (Leyton-Brown and Shoham, 2013). It attempts to fashion rules with specific properties to produce a certain outcome. The following sections introduce the relevant concepts and notations used in mechanism design theory (Börgers, 2015; Mas-Colell et al., 1995; Fudenberg and Tirole, 1993).

2.1.1 Bayesian mechanisms

Bayesian mechanism design studies the implementation of a mechanism in Bayesian Nash equilibrium. In *Bayesian* games, the information about the preferences of the agents is private to these agents and are described by given probability distributions (Korb and Nicholson, 2010). Bayesian mechanisms induce games of incomplete information. Specific social choice functions use these the agents' privately known preferences to determine an allocation that satisfies a number of desirable properties. In particular, games induced by mechanisms are analyzed for Bayesian Nash equilibria. One prominent result in mechanism design theory – the revelation principle – provides strong implications on mechanism implementations in Bayesian Nash equilibrium.

2.1.1.1 Private information

When some agents do not know the preferences of the others, the game is said to have *private information*. Bayesian mechanism design assumes that this private information follows a set of publicly known probability distributions. This assumption is referred to as the *independent private values assumption* (Milgrom and Weber, 1982). In models with independent private values, the private information of each agent is independent from that of other agents and follows commonly known probability distributions. An agent's private information is also known as the agent's *type* in mechanism design theory. Thus, each agent is defined by his type. Let $\theta_i \in \Theta_i$ denote the type of agent a_i, where Θ_i is the associated type space. Let F_i denote the publicly accessible probability distribution on Θ_i. In game theory, such settings are modeled by games of *incomplete information*, which were first introduced by Harsanyi (1967-1968).

Definition 2.1. Following Nisan et al. (2007), a *game of incomplete information* with $I \in \mathbb{N}$ agents is characterized by:

1. For each agent a_i, there exists a set of strategies S_i. The strategy of a_i is a function $s_i : \Theta_i \to S_i$, which determines a_i's action based on his type $\theta_i \in \Theta_i$.

2. For each agent a_i, there exists a set of types Θ_i, as well as the publicly accessible probability distribution F_i on Θ_i. The value $\theta_i \in \Theta_i$ is the type (i.e., the private information) of a_i. The quantity $F_i(\theta_i)$ is the probability realization that a_i is of type θ_i.

3. For each agent a_i, there exists a utility function $u_i : S_1 \times \cdots \times S_I \times \Theta_i \to \mathbb{R}$. Agents are assumed to be *risk neutral*. Game theory denotes such utility functions as *quasi-linear* because they exhibit linear and separable dependence on money.

In games of incomplete information, each agent must select his strategy based on his own type θ_i, not knowing the other agents' types $\theta_{-i} = (\theta_1, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots, \theta_I)$. The agent only knows the common probability distributions of the other agents' types. The main goal of mechanism design theory is to implement rules such that these strategies are in equilibrium. To describe the implementation of these rules, functions are required that translate the agents' private information into a collective choice. These functions are referred to as *social choice functions*.

2.1.1.2 Social choice functions

Since the preferences of the agents depend on their privately observed types $\{\theta_1, \ldots, \theta_I\}$, the collective decision itself must be a function of these types. This social choice function takes as input the types of all the agents and determines an allocation in Ω, where Ω denotes the set of all possible allocation alternatives.

Definition 2.2. A *social choice function* is a function $f_{sc} : \Theta_1 \times \cdots \times \Theta_I \to \Omega$, which maps the types $(\theta_1, \ldots, \theta_I) \in \Theta_1 \times \cdots \times \Theta_I$ of all agents to an allocation in Ω.

A mechanism defines a set of rules that govern the actions of the agents in games of incomplete information. The rules specify how these actions translate into a collective allocation outcome. Hence, a mechanism *indirectly* implements a social choice function by providing a platform on which the agents interact to make a collective choice.

Definition 2.3. A *mechanism* is a tuple $\mathscr{M} = (S_1, \ldots, S_I, g(\cdot))$, where $S_1, \ldots, S_I$ are the strategy sets of the agents and $g : S_1 \times \cdots \times S_I \to \Omega$ is an outcome function.

A mechanism induces a game of incomplete information. The outcome function g specifies the rules for how the actions of the agents translate into a social choice. The following example illustrates these definitions applied to first-price auctions.

Example 2.1. Consider a first-price sealed-bid auction. This auction is a mechanism where I agents compete for a single, indivisible good. The agent proposing the highest bid receives the good and transfers a payment equal to this bid. Formally, the strategy set of each agent is $S_i = \mathbb{R}^+$, the set of all positive real numbers, and, given the individual bids $b_1, \ldots, b_I$ of each agent, the outcome function of the mechanism is specified by $g(b_1, \ldots, b_I) = (y_1(b_1, \ldots, b_I), \ldots, y_I(b_1, \ldots, b_I), t_1(b_1, \ldots, b_I), \ldots, t_I(b_1, \ldots, b_I))$, where

$$y_i(b_1, \ldots, b_I) = \begin{cases} 1, & \text{if and only if } i = \min\{j : b_j = \max\{(b_1, \ldots, b_I)\}\} \\ 0, & \text{otherwise} \end{cases}$$

denotes the decision variable, and

$$t_i(b_1, \ldots, b_I) = -b_i y_i(b_1, \ldots, b_I)$$

defines the payment transferred by the winner.

Because a mechanism induces a game of incomplete information, research in mechanism design theory is concerned with implementing the social choice function $f_{sc}(\cdot)$ such that there is an equilibrium of the game that yields the same outcome as $f_{sc}(\cdot)$ for every possible type vector $(\theta_1, \ldots, \theta_I) \in \Theta_1 \times \cdots \times \Theta_I$. The standard solution concept for games of incomplete information is that of a Bayesian Nash equilibrium (Fudenberg and Tirole, 1993). The following section characterizes the concept of Bayesian Nash equilibria.

2.1.1.3 Bayesian Nash equilibria

When searching for equilibria in a game of incomplete information, two different solution concepts can be applied to predict the agents' strategic behavior: *dominant strategies* and *Bayesian Nash equilibria*. In a dominant strategy mechanism, each agent's optimal type announcement is independent of all other agents' announcements; that is, there exist no publicly known probability distributions whatsoever. In applications of dominant strategies, the agents emanate from the "worst-case" behavior of the other agents. In contrast, implementing social choice functions in the Bayesian context depends on the realization of the agents' *expected* utilities. Since $g : S_1 \times \cdots \times S_I \to \Omega$ is the outcome function, the strategies of the agents translate into an allocation alternative $x \in \Omega$. Each agent is assumed to be an expected utility maximizer, whose quasi-linear utility function is $u_i(x, \theta_i)$, given his type θ_i and an allocation alternative $x \in \Omega$. The following definition characterizes Bayesian Nash equilibria, initially introduced by the seminal work of Harsanyi (1967-1968).

Definition 2.4. The strategy combination $s(\cdot) = (s_1(\cdot), \ldots, s_I(\cdot))$ is a *Bayesian Nash equilibrium* of mechanism $\mathscr{M} = (S_1, \ldots, S_I, g(\cdot))$ if, for all $\theta_i \in \Theta_i$ and all $\hat{s}_i \in S_i$,

$$\mathbb{E}_{\theta_{-i}}[u_i(g(s_i(\theta_i), s_{-i}(\theta_{-i})), \theta_i)] \geq \mathbb{E}_{\theta_{-i}}[u_i(g(\hat{s}_i, s_{-i}(\theta_{-i})), \theta_i)]. \tag{2.1}$$

Hence, the strategy profile $s_1, s_2, \ldots, s_I$ is a Bayesian Nash equilibrium if, for every agent a_i and every type θ_i, strategy $s_i(\theta_i)$ is the best response of a_i to the strategies $s_{-i}(\theta_{-i})$ of all other agents, when a_i is of type θ_i, in expectation over the types of all other agents $\mathbb{E}_{\theta_{-i}}[\cdot]$. The following definition characterizes the implementation of a mechanism in the Bayesian sense.

Definition 2.5. The mechanism $\mathscr{M} = (S_1, \ldots, S_I, g(\cdot))$ *implements the social choice function* $f_{sc}(\cdot)$ *in Bayesian Nash equilibrium* if there exists a Bayesian Nash equilibrium $(s_1(\cdot), \ldots, s_I(\cdot))$ such that $g(s_1(\theta_1), \ldots, s_I(\theta_I)) = f_{sc}(\theta_1, \ldots, \theta_I)$ for all types $(\theta_1, \ldots, \theta_I) \in \Theta_1 \times \cdots \times \Theta_I$.

2.1.1.4 The revelation principle

In general, there may be a large set of possible mechanisms that implement a social choice function. However, the *revelation principle* states that, in searching for mechanisms that implement a social choice function, attention can be restricted to one class of very simple mechanisms: the *direct revelation mechanisms*. In a direct revelation mechanism, all agents are induced to directly disclose their types. Then, given the announced type vector $(\hat{\theta}_1, \ldots, \hat{\theta}_I)$, the mechanism determines an allocation such that $x = f_{sc}(\hat{\theta}_1, \ldots, \hat{\theta}_I)$.

Definition 2.6. A *direct revelation mechanism* is a mechanism in which $S_i = \Theta_i$ for each agent a_i and $g(\theta_1, \ldots, \theta_I) = f_{sc}(\theta_1, \ldots, \theta_I)$ for all types $(\theta_1, \ldots, \theta_I) \in \Theta_1 \times \cdots \times \Theta_I$.

A direct revelation mechanism mandates all agents to directly announce their types to the mechanism. Then, based on the reported types, the mechanism computes the strategies that each agent would have used in the original mechanism. Finally, the mechanism implements the outcome function $g(\cdot)$ prescribed in the given game for these strategies. Myerson and Satterthwaite (1983, p. 267-268) informally describe the main idea of the revelation principle as follows:

> *The essential idea is that, given any equilibrium of any bargaining game, we can construct an equivalent incentive-compatible direct mechanism by first asking the buyer and seller each to confidentially report his valuation, then computing what each would have done in the given equilibrium strategies with these valuations, and then implementing the outcome (transfer of money and object) as in the given game for this computed behavior. If either individual had any incentive to lie to us in this direct mechanism, then he would have had an incentive to lie to himself in the original game, which is a contradiction of the premise that he was in equilibrium in the original game.*

The following proposition formalizes the revelation principle in the Bayesian context.

Proposition 2.1 (Bayesian revelation principle)**.** *If there exists a mechanism* $\mathscr{M} = (S_1, \ldots, S_I, g(\cdot))$ *that implements the social choice function* $f_{sc}(\cdot)$ *in Bayesian Nash equilibrium, then* $f_{sc}(\cdot)$ *is Bayesian incentive compatible.*

Proofs of the proposition can be found in the work of Mas-Colell et al. (1995) as well as that of Fudenberg and Tirole (1993) and are omitted here.

Early versions of the revelation principle were developed by the seminal work of Gibbard (1973) and further extended by the work of Dasgupta et al. (1979), Myerson (1979, 1981), Harris and Townsend (1981), and Harris and Raviv (1981). The revelation principle greatly simplifies the search for mechanisms that implement social choice functions. Without loss of generality, the mechanism designer can restrict attention to direct revelation mechanisms, in which all agents simply announce their types and, based on these announcements, the mechanism's outcome function equals the social choice function. The set of equilibrium outcomes induced by direct revelation mechanisms are those mechanisms that satisfy incentive compatibility. Section 2.1.2 further explores the notion of *incentive compatibility*.

2.1.2 Desirable mechanism properties

The design objective of a mechanism is to arrive at an efficient allocation with respect to a set of economically desirable properties. These properties include (i) incentive compatibility because agents are assumed to act strategically, (ii) individual rationality because participation in the mechanism must be voluntary, (iii) budget balance because no outside party exists to subsidize the mechanism, and (iv) social welfare maximizing. This section introduces these properties in the Bayesian context.

2.1.2.1 Incentive compatibility

To determine the optimal equilibrium outcome from all possible mechanisms, this research adopts the standard approach by invoking the revelation principle. As discussed in Section 2.1.1.4, attention can be restricted to incentive compatible direct mechanisms without loss of generality. Thus, if a social choice function is implementable in Bayesian Nash equilibrium, it is also Bayesian incentive compatible according to the following definition.

Definition 2.7. The social choice function $f_{sc}(\cdot)$ is *Bayesian incentive compatible* if, for all types $\theta_i \in \Theta_i$, the strategy combination $s(\cdot) = (\theta_1, \ldots, \theta_I)$ is a Bayesian Nash equilibrium of the direct revelation mechanism $\mathscr{M} = (\Theta_1, \ldots, \Theta_I, f_{sc}(\cdot))$. Thus, for all $\hat{\theta}_i \in \Theta_i$,

$$\mathbb{E}_{\theta_{-i}}[u_i(f_{sc}(\theta_i, \theta_{-i}), \theta_i)] \geq \mathbb{E}_{\theta_{-i}}[u_i(f_{sc}(\hat{\theta}_i, \theta_{-i}), \theta_i)]. \tag{2.2}$$

2.1.2.2 Individual rationality

In games of incomplete information, the utilities of the agents are *expected utilities*. The expectation of these utility functions crucially depends on the *point in time* when the types of the agents are observed by the agents. The literature provides three evaluation stages to describe the appropriate timing (Holmström and Myerson, 1983): *ex post*, *interim*, and *ex ante*. At each stage, there are different expectations for the agent's utility function.

In the *ex post* timing stage of the direct revelation mechanism, all types of the agents are publicly known to all agents. At this point in time, all types have been reported to the mechanism and an allocation $x \in \Omega$ has been selected. In this case, agent a_i's *ex post* utility function is given by $u_i(x, \theta_i)$, when a_i is of type θ_i. That is, it is not necessary to take expectations because all type are publicly known.

In the *interim* timing stage of the direct revelation mechanism, each agent privately observes his own type but does not know any other agents' types. In fact, the other agents' types are given by publicly known probability distributions. At this point in time, the agents have not yet reported their types. In this case, agent a_i's *interim* utility function evaluates to $\mathbb{E}_{\theta_{-i}}[u_i(x, \theta_i)]$, when a_i is of type θ_i. Function $\mathbb{E}_{\theta_{-i}}$ takes expectations with respect to all other agents' types θ_{-i} conditional on a_i's known type θ_i.

In the *ex ante* timing stage of the direct revelation mechanism, the agents know neither their own type nor the types of any other agents. The agents have not observed any type but instead rely on the publicly accessible probability distributions. In this case, agent a_i's *ex ante* utility function is $\mathbb{E}[u_i(x, \theta_i)]$, when a_i is of type θ_i. Function $\mathbb{E}$ denotes the expectation with respect to all agents' types.

These three evaluation stages play a central role in defining the mechanism's Bayesian incentive compatibility constraint (2.2). This constraint, however, is not the only desirable property of the mechanism. The agent participates in the mechanism only if his expected utility is non-negative. If participation in the mechanism is voluntary, agent a_i's decision whether to participate must be *individually rational*. Hence, imposing an individual rationality constraint on the mechanism is desirable but may limit the set of social choice functions that can successfully be implemented.

Suppose that agent a_i of type θ_i receives a utility of $u_i(x, \theta_i) = 0$ if he withdraws from the mechanism with allocation outcome $x = f_{sc}(\theta_i, \theta_{-i})$. This utility is also referred to as the *outside option* of that agent in mechanism design theory. To ensure a_i's participation in the mechanism, the individual rationality constraint to be imposed depends on the current evaluation stage. An agent participates in the mechanism only if a non-negative expected utility can be guaranteed. In the ex post stage, this condition means $u_i(x, \theta_i) \geq 0$; thus, the agent can withdraw from the mechanism even *after* all types are publicly known and the final allocation $x \in \Omega$ is selected. However, if agent a_i is only able to withdraw from the mechanism *after* all agents have learned their types but *before* they report them, then agent a_i's *interim* individual rationality constraint is given by $\mathbb{E}_{\theta_{-i}}[u_i(x, \theta_i)] \geq 0$. Finally, if agent a_i can withdraw only *before* all agents observe their types, his *ex ante* individual rationality criterion is $\mathbb{E}[u_i(x, \theta_i)] \geq 0$.

2.1.2.3 Budget balance

As discussed in Example 2.1, allocation mechanisms prescribe monetary payments that are paid to or received by the agents for taking participation in an allocation. In general, the mechanism must omit any independent intermediary and facilitate distributed decision-making among the agents (Egri and Váncza, 2013). This requirement implies that all payments must be distributed among the agents, and the mechanism must raise sufficient revenue from the agents to cover the cost for running the platform. Game theory denotes this property as *budget balance* (Fudenberg and Tirole, 1993). In a budget-balanced mechanism, all monetary transfers between the agents must sum to zero. If this condition does not hold, the mechanism must be subsidized continuously by a third party. Similar to the individual rationality constraint, the budget balance constraint can be imposed at different timing stages (cf. Section 2.1.2.2). This research focuses on the ex ante budget balance constraint.

Definition 2.8. The social choice function $f_{sc} : \Theta_1 \times \cdots \times \Theta_I \to \Omega$ satisfies *ex ante budget balance* if the expectation of all monetary transfers

$t_1(\theta_1,\ldots,\theta_I),\ldots,t_I(\theta_1,\ldots,\theta_I)$ made by the agents sums to zero. Thus, for all $(\theta_1,\ldots,\theta_I)\in\Theta_1\times\cdots\times\Theta_I$,

$$\mathbb{E}\Big[\sum_{i=1}^{I} t_i(\theta_1,\ldots,\theta_I)\Big]=0. \tag{2.3}$$

2.1.2.4 Optimality

The ultimate goal of an allocation mechanism is to satisfy Pareto optimality (Sen, 1970). In a Pareto optimal allocation, it is not possible to increase one agent's utility without decreasing the utility of at least one other agent. Pareto optimality and ex post optimality are often used interchangeably in mechanism design theory.

Definition 2.9. The social choice function $f_{sc}:\Theta_1\times\cdots\times\Theta_I\to\Omega$ is *ex post optimal* if, for no type profile $(\theta_1,\ldots,\theta_I)$, there is an allocation alternative $x\in\Omega$ such that $u_i(x,\theta_i)\geq u_i(f_{sc}(\theta_i,\theta_{-i}),\theta_i)$ for every i, and $u_i(x,\theta_i)>u_i(f_{sc}(\theta_i,\theta_{-i}),\theta_i)$ for some i.

Ex post optimal mechanisms are referred to as *first-best* mechanisms in the literature (Börgers, 2015). Because agents use quasi-linear utility functions, a social choice function is ex post optimal if it maximizes the sum over all utilities. The sum over all agents' utilities is also known as the *utilitarian* social welfare (Chevaleyre et al., 2006). Thus, the ultimate design objective is to solve

$$\max_{x\in\Omega}\sum_{i=1}^{I} u_i(x,\theta_i), \tag{2.4}$$

subject to incentive compatibility, individual rationality, and budget balance. However, standard impossibility theorems from mechanism design theory assert that there is no ex post optimal social choice function that satisfies these requirements simultaneously (Myerson and Satterthwaite, 1983). Since ex post optimality is unattainable, research in mechanism design proposes alternative objective functions. Myerson and Satterthwaite (1983) seek mechanisms that maximize the *ex ante* (expected) social welfare, subject to incentive compatibility and individual rationality. In mechanism design theory, such mechanisms are also known as *ex ante efficient* (Rustichini et al., 1994).

Definition 2.10. The social choice function $f_{sc} : \Theta_1 \times \cdots \times \Theta_I \to \Omega$ is *ex ante optimal* if, for no type profile $(\theta_1, \ldots, \theta_I)$, there is an allocation alternative $x \in \Omega$ such that $\mathbb{E}[u_i(x, \theta_i)] \geq \mathbb{E}[u_i(f_{sc}(\theta_i, \theta_{-i}), \theta_i)]$ for every i, and $\mathbb{E}[u_i(x, \theta_i)] > \mathbb{E}[u_i(f_{sc}(\theta_i, \theta_{-i}), \theta_i)]$ for some i.

Unlike ex post optimality, ex ante optimality only requires the maximization of the welfare's *expectation* across all agent types. This specific objective function places equal welfare weights on the utilities of the agents. Mechanisms that are ex ante optimal are known as *second-best* mechanisms (Börgers, 2015).

Definition 2.11. A *second-best* mechanism is one that maximizes the expected social welfare among all incentive compatible and individually rational mechanisms.

Thus, the optimization problem of a second-best mechanism is the maximization of the expected social welfare, which is given by

$$\max_{x \in \Omega} \mathbb{E}\Big[\sum_{i=1}^{I} u_i(x, \theta_i)\Big], \tag{2.5}$$

subject to incentive compatibility, individual rationality, and budget balance. Certainly, second-best mechanisms are not socially optimal because ex post optimality cannot be achieved. However, the outcome of a second-best mechanism can be compared to the outcome of the associated first-best mechanism if it were to exist. To compute the outcome of the associated first-best mechanism, one would assume that all information is publicly given. Then, in the presence of public information, any incentive compatibility constraints imposed by the mechanism are dropped because the traders cannot misrepresent their types. Using the first-best outcome as a benchmark, one can then measure the efficiency of the second-best mechanism. In particular, the *efficiency* of the mechanism is defined as the proportion of the second-best outcome relative to the outcome that the first-best mechanism would achieve had it existed.

Definition 2.12. Let V denote the outcome of a second-best mechanism and let V^* denote the (socially optimal) outcome of the associated first-best mechanism if it were to exist. Then the *efficiency* of the second-best mechanism is defined by the ratio $\frac{V}{V^*}$. The associated *asymptotic efficiency* is defined as the efficiency of the mechanism for a market, where the number of participating agents approaches infinity (Gresik and Satterthwaite, 1989; Rustichini et al., 1994).

With these basic definitions from mechanism design theory in place, a set of requirements on allocation mechanisms for electronic services can be developed. In the following section, these requirements are specified in detail.

2.2 Requirements on allocation mechanisms for electronic services

This section characterizes the set of economic requirements on mechanisms that must be fulfilled in order to solve the allocation problem. The first three requirements (double-sided competition, quality awareness, and private information) are due to the characteristics of electronic services and markets. The last four requirements (incentive compatibility, individual rationality, budget balance, and optimality) emerge from more general characteristics of mechanism design theory (cf. Section 2.1.2).

2.2.1 Double-sided competition

Electronic services are offered by *multiple* providers to *multiple* customers on electronic marketplaces (e.g., Avnet Cloud Marketplace, 2016). Such marketplaces are platforms that bring together two distinct groups of users, namely service providers and service customers (Eisenmann et al., 2006). Value increases as the marketplace matches demand from both sides. The main challenge of such double-sided marketplaces is to choose an appropriate pricing strategy because one side must incorporate the effects from the other market side's growth and willingness to pay for the services. On double-sided marketplaces, these effects may be either positive or negative. For instance, increasing the number of customers may be more or less advantageous for the customers' side of the market (i.e., same-side effect) or for the providers' side of the market (i.e., cross-side effect). Therefore, due to the growing number of customers and providers on both market sides, the allocation mechanism for electronic services must support competition on double-sided marketplaces.

Requirement 1 (Double-sided competition)**.** *The allocation mechanism must support double-sided competition between multiple customers and multiple providers.*

2.2.2 Quality awareness

The growth of information and communication technology has significantly influenced the way electronic services are developed and provided. Even though the concept of electronic services has been discussed extensively in research literature, there is no agreement regarding a common definition (Gadrey, 2000; Lindgrena and Jansson, 2013). In particular, business services that use electronic networks as an infrastructure are referred to as *electronic services*. Here, the concept of a "service" is restricted to *software services* (Papazoglou, 2008); that is, electronic services are software services provided over electronic networks (Rust and Kannan, 2002).

Customers of electronic services rely on service providers to supply the appropriate computing needs (Buyya et al., 2009). They expect a certain quality from providers, control in transactions, and a certain level of variety in the choice of service settings. Such quality attributes represent customer expectations with respect to the demanded electronic service. These qualitative service characteristics are referred to as *quality of service* (QoS). The QoS belongs to the class of *non-functional properties* of an electronic service and describes how the service performs (O'Sullivan et al., 2002). These non-functional service properties include measurable qualitative characteristics such as availability and throughput. In contrast to a service's non-functional properties, its *functional properties* specify what the service actually does; that is, the actions the service performs as well as the valid inputs and outputs of the service. Electronic services that exhibit identical functional properties are denoted *standardized* electronic services. The OASIS consortium has recently promoted the industrial standardization of electronic services in Cloud computing settings (OASIS, 2016).

This research focuses on *non-functional* properties associated with standardized electronic services. In particular, each service is characterized by its quality. Typically, this QoS comprises a set of quantifiable and measurable QoS parameters. In this research, each electronic service is described by a single QoS parameter, which determines the qualitative characteristics of the service. Thus, QoS is modeled as a one-dimensional parameter.

Requirement 2 (Quality awareness)**.** *The allocation mechanism must internalize the QoS desired by customers and the QoS offered by providers.*

2.2.3 Private information

The exact realization of the QoS desired by a customer is private to the customer and cannot be observed by any other agent. Similarly, the actual QoS offered by a provider is private to that provider and remains unknown to other agents (Bhargava and Sun, 2008). For instance, customers may have imperfect knowledge about QoS actually delivered by providers and cannot discern low and high quality providers. On the other hand, providers might be unaware of the current QoS needs of customers. Such mutual uncertainty concerning QoS entails negative business consequences (Bazerman and Gillespie, 1999). However, service providers can exploit QoS uncertainties to induce customers to choose the appropriate QoS. In such settings, high-quality customers pay a higher price for the electronic service. At the same time, service providers use differentiated QoS levels to distinguish themselves from their competitors and to create a higher value for the electronic service (Papazoglou, 2008). Therefore, the QoS demanded by customers and the QoS offered by providers must be represented as private information in the mechanism.

Requirement 3 (Private information)**.** *The allocation mechanism must account for private information about the QoS desired by customers and the QoS offered by providers.*

2.2.4 Incentive compatibility

When strategic agents negotiate over electronic services, the allocation mechanism must provide adequate incentives to motivate truthful behavior among the agents. In settings with private information about the QoS, an allocation mechanism is Bayesian incentive compatible if honest type declaration is a Bayesian Nash equilibrium for all customers and providers (cf. Section 2.1.2.1).

Requirement 4 (Incentive compatibility)**.** *The allocation mechanism must satisfy incentive compatibility for all customers and providers.*

2.2.5 Individual rationality

Customers and providers only participate in the mechanism if they derive non-negative utility from participation. They refuse to participate in any allocation mechanism if their expected utility is negative. Therefore, the allocation mechanism must ensure that participation in the mechanism is voluntary to that agent (cf. Section 2.1.2.2).

Requirement 5 (Individual rationality)**.** *The allocation mechanism must satisfy individual rationality for all customers and providers.*

2.2.6 Budget balance

On markets for trading electronic services, no external source of funds exists that is willing to permanently subsidize the platform. Instead, the underlying allocation mechanism must facilitate distributed decision-making by dispensing all monetary transfers among the agents. Thus, payments and compensations must offset each other. The allocation mechanism satisfies budget balance if all monetary transfers sum to zero for all agents (cf. Section 2.1.2.3). If a mechanism does not fulfill budget balance, it requires an external net source of funds to continually subsidize the allocation mechanism. A budget-balanced mechanism does not run a permanent monetary deficit but is independent of any external source of funds. Therefore, all allocation mechanisms require budget balance.

Requirement 6 (Budget balance)**.** *The allocation mechanism must satisfy budget balance.*

2.2.7 Optimality

Ex post optimality cannot be achieved if incentive compatibility, individual rationality, and budget balance are required (cf. Section 2.1.2.4). Because private information is present, this research derives a second-best mechanism, which maximizes the expected social welfare, subject to incentive compatibility, individual rationality, and budget balance. Hence, the objective function of the mechanism is the maximization of the expected social welfare.

Requirement 7 (Optimality). *The allocation mechanism must be ex ante optimal; that is, it must maximize the expected social welfare.*

On the basis of these requirements, the following section presents the analysis of pertinent mechanisms in the literature that are closely akin to the work in this thesis. This analysis lays the foundation for the design of any allocation mechanism that attempts to satisfy the requirements elaborated in this section.

2.3 Analysis of allocation mechanisms for electronic services

This section analyzes extant allocation mechanisms for electronic services that are closely related to the mechanism proposed in this thesis. These mechanisms are divided into five different groups to provide a fundamental demarcation between the requirements derived in Section 2.2. Each group is defined by pairing together the salient mechanism requirements of double-sided competition, quality awareness, incentive compatibility, and optimality. Within each group, every allocation mechanism is analyzed in four steps. First, the formal model of the mechanism is described. Second, the findings reported for the mechanism are summarized. Third, the assumptions that underlie the model and its limitations are identified. Finally, the applicability and the relevance for this thesis are discussed.

2.3.1 Single-sided, quality-aware

This section provides an analysis of three single-sided mechanisms for allocating electronic services that maintain quality awareness. While the first two approaches by Che (1993) and Bichler and Kalagnanam (2005) consider multidimensional auctions from the perspective of a profit-maximizing intermediary, Blau et al. (2010) focus on quality-aware mechanisms that maximize the social welfare.

Design competition through multidimensional auctions

Among the first auction models to integrate both price and quality is the multidimensional auction model proposed by Che (1993). This model is based on prior research in information economics (Dasgupta and Spulber, 1990). A single customer announces a publicly known scoring rule to multiple providers who compete to win the auction. Price and quality preferences are aggregated in the utility function of each agent. Che takes the standard approach in mechanism design theory by invoking the revelation principle to identify the optimal mechanism. Using the associated optimal outcome as a benchmark, Che studies the indirect mechanism implementation through three distinct auction formats. In the first-score auction, each provider submits a bid and, upon winning, delivers the offered quality at the offered price. In the second-score auction, the winner must match the second-highest score in the auction. In the second-preferred-offer auction, the winner must match the exact quality-price combination of the second-highest score. Che studies the performance of these different auction formats in favor of the profit-maximizing customer. Subsequently, he analyzes the optimal quality choice of the winning provider under alternative auction formats. Branco (1997) extends Che's work by integrating correlated cost types of the providers into the auction mechanism.

Che (1993) finds that the performance of the three different auction formats critically depends on the customer's design of the scoring rule and his commitment power. If the customer chooses a scoring rule that reflects his true utility (i.e., the naive scoring rule), then all three auction formats result in the same expected utility to the customer. Che also finds that the naive scoring rule leads to excessive quality under first- and second-score auctions. This result implies that the customer has an incentive to deviate from the naive scoring rule. The magnitude of this deviation corresponds to the quality distortion identified by the revelation principle. Consequently, Che derives the optimal scoring rule that counterbalances the excessive quality appearing in the naive scoring rule. Hence, the optimal scoring rule systematically discriminates against quality. Finally, Che finds that the first- and second-score auction can implement the optimal outcome, while the second-preferred auction cannot.

The study of Che (1993) is based on the following assumptions. Although Che identifies the optimal mechanism for allocating objects of distinct quality levels, the mechanism assumes a single customer on the market. It remains unclear to what extent the optimal scoring rule must be adapted once multiple customers enter the market. Moreover, private information exists only on the providers' side of the market. For this reason, each provider's cost is a function of the private cost information and the offered quality parameter. How the existence of private information on the customer side of the market affects the optimal mechanism definition cannot be assessed by the model. Finally, Che's model is not instantiated in an experimental evaluation.

Although the study of Che (1993) assumes single-sided markets, it presents one of the first auctions to make the quality of the traded object an intrinsic part of the mechanism. Since then, much of the succeeding research has built upon scoring rules for aggregating multiple attributes such as price and quality (Asker and Cantillon, 2008). With respect to this thesis, it is crucial that Che invokes the revelation principle for identifying the optimal scoring rule (cf. Section 2.1.1.4). As a result, Che finds that in the optimal revelation mechanism, quality is distorted downward to limit the information rents of providers. This finding is important for understanding the strategic behavior of the agents in this thesis. Here, the definition of a provider's virtual cost captures the fact that providers have an incentive to overstate their true cost in order to raise the transaction price. Because Che assumes single-sided competition, however, the influence that customers can exert on the price cannot be assessed. Thus, this thesis drops the assumption of single-sided competition by defining a virtual valuation that reflects the strategic behavior of customers (cf. Section 4.1.1).

Configurable offers and winner determination in multi-attribute auctions

Specific auctions for the procurement of electronic services with multiple attributes have also been studied. Instead of awarding the winning contracts to a single provider (as in the model by Che (1993)), Bichler and Kalagnanam (2005) extend the so-called *sole sourcing* auction format to the case of *multiple sourcing*. Multiple sourcing auction formats support the allocation of the customer's demand to multiple providers. On the other hand, Bichler and Kalagnanam study the configurability of the bids in multi-attribute auctions. Configurable offers enable providers to specify multiple QoS attributes and price information for each attribute. Bichler and Kalagnanam focus on winner determination problems that arise in multi-attribute auctions with multiple sourcing and configurable offers. For describing these winner determination problems, they define scoring functions based on the preferences of a customer. Thus, each possible configuration of a provider's configurable offer is reflected in a function of price per quantity and QoS attributes. The pricing scheme of the allocation mechanism is based on additive pricing functions, where the price dependence on an attribute is specified as a markup over a fixed base price. In a set of numerical simulations, Bichler and Kalagnanam study the complexity properties and computational issues of the different winner determination problems, which arise from the associated allocation setting with configurable offers. Their model is instantiated in a numerical simulation by executing the artifact with artificial data. In fact, an implementation of the associated Multidimensional Auction Platform (MAP) has been used in a large-scale procurement marketplace for the retail industry (Bichler et al., 2002).

Bichler and Kalagnanam (2005) find that the winner determination problem can be solved with typical attribute configurations in the order of a few seconds by using a branch-and-bound approach. However, imposing additional constraints on the minimum number of winning providers results in an exponential runtime for the allocation problem. In particular, they find that the allocation problem of configurable offers is considerably easier to solve without homogeneity constraints. Homogeneity constraints can be imposed on the winner determination problem to enforce homogeneity of a certain QoS attribute in the set of winning bids. Their findings indicate that if homogeneity constraints are not considered, the overall winner determination problem can be divided in a set of smaller problems, in which the optimal configuration for each configurable offer is selected based on a customer's scoring function. Again, Bichler and Kalagnanam find that imposing homogeneity constraints for configurable offers leads to computational infeasibility.

The approach of Bichler and Kalagnanam (2005) is subject to the following assumptions. While Bichler and Kalagnanam focus on formulating the winner determination for their allocation problem, they assume a scoring rule that is fixed prior to the allocation process. Taking the scoring rule as given by the customer, however, does not allow for assessing the optimality of the allocation outcome. In particular, their formal model does not provide sufficient constructs to analyze incentive compatibility or individual rationality. Thus, it remains unclear to what extent the suggested pricing rules entail truthful behavior of the agents. Moreover, similar to Che (1993), the model of Bichler and Kalagnanam is limited to a single customer, which is why their approach cannot be used to examine double-sided competition.

Although the mechanism studied by Bichler and Kalagnanam assumes payment schemes for which incentive compatibility and individual rationality are not assessed, their mechanism is relevant to the mechanism developed in this thesis because Bichler and Kalagnanam formulate the allocation problem as a constrained optimization problem. Constrained optimization for multi-attribute auctions has been used in subsequent research to study clearing algorithms in double auctions (Engel et al., 2006) as well as the preference-based selection of configurable web services (Lamparter et al., 2007). Both studies successfully model the winner determination problem as an optimization problem subject to a set of multi-attribute feasibility constraints. This thesis embraces this approach by representing the mechanism's optimality criterion as a constrained optimization problem. However, in order to facilitate payments that satisfy incentive compatibility and individual rationality, the optimization problem must be extended by additional constraints. These constraints must not only ensure the feasibility of the optimization problem (as in the work of Bichler and Kalagnanam). In fact, they must also guarantee that the payment schemes satisfy incentive compatibility and individual rationality. Therefore, the mechanism developed in this thesis drops the assumption of non-existing constraints for incentive compatibility and individual rationality by imposing additional constraints that ensure incentive compatible and individually rational payments (cf. Section 3.6).

A multidimensional procurement auction for trading composite services

A particular scenario of electronic service allocation is the allocation of services that are part of a composite service. A composite service is the assembly and invocation of multiple pre-existing, standardized services, possibly offered by diverse providers to complete the functionality of a multi-step business process (Papazoglou, 2008). For this reason, research in multidimensional procurement auctions is concerned with trading composite electronic services (Blau et al., 2009, 2010). In contrast to providing traditional service bundles, the provision of a complex composite service highly depends on the accurate execution order of each service component and only generates value for the customer if these components occur in a valid sequence. Traditional negotiation mechanisms such as combinatorial auctions and their appropriate protocol implementation are not sufficient to reach the desired service allocation. Thus, Blau et al. (2010) study a mechanism that forms the demanded complex service through the allocation and pricing of individual service components. Their model is informed by prior research in algorithmic mechanism design (Nisan and Ronen, 2001). For a composite service, the valuation of each service requester highly depends on the accurate sequence of its functional parts. Hence, these composite services only generate value through a valid order of their components. To capture the preferences of customers regarding the different QoS, Blau et al. (2010)

draw from the work of Asker and Cantillon (2008) and use scoring functions to map multiple attributes to a single scoring value. The allocation of service components is computed based on the scoring value for QoS demands of the customer. The proposed mechanism is individually rational for service providers, and the resulting allocation maximizes the social welfare across all agents. The proposal supports incentive compatibility with respect to QoS characteristics and prices in weakly dominant-strategy implementations for all service providers. Using an experimental evaluation, Blau et al. (2010) study a set of alternative strategies service providers might pursue by forming strategic coalitions among the providers.

Blau et al. (2010) find that bundling strategies due to coalition formation lead to a significantly higher expected payoff for less competitive service providers. These providers offer their services at high prices or at QoS levels that are not sufficiently valuable for customers. In this way, such providers can combine their offers with more attractive components offered by other providers in order to increase their chance of being allocated and the expected payoff. In contrast, competitive service providers tend to form coalitions with less competitive service providers. Blau et al. find that customers might view their offers as less favorable although the provision costs of providers can be reduced by coalition formation. This risk results in a reduced likelihood of being allocated.

The study of Blau et al. (2010) makes the following assumptions for the proposed mechanism. Although Blau et al. design an allocation mechanism for electronic services with QoS attributes from the perspective of a social planner (i.e., a social welfare maximizer), the mechanism does not satisfy budget balance. Thus, a third party must permanently subsidize the mechanism in order to maintain its operation. Further, the mechanism assumes a single customer on the market, who is demanding services from multiple service providers. Thus, competition can be realized among service providers alone.

Although the allocation mechanism presented by Blau et al. (2010) assumes a market with a single customer, it is relevant to the approach taken in this thesis because Blau et al. focus on maximizing the social welfare instead of the customer's profit (as in Bichler and Kalagnanam (2005)). Hence, they introduce a social choice by allocating the demanded services from multiple self-interested providers. Subsequent mechanisms for allocating electronic services embrace this concept by representing the objective function of the mechanism in terms of a welfare maximization function (Haas et al., 2013). This thesis adopts this approach by defining the mechanism's objective function as the sum of all agents' utilities (cf. Section 3.6). In Blau et al. (2010), each provider's internal cost structure and QoS is known only to that provider. For this reason, the types of the providers are used to internalize these private costs and QoS into the allocation mechanism. In order to support multiple competing customers, however, the model presented in this thesis must not only integrate the providers' private information concerning their offered QoS. Instead, it is necessary to represent the QoS desired by customers as private information also, thus making the privately known QoS an intrinsic part of the mechanism (cf. Section 3.2).

2.3.2 Single-sided, approximately socially optimal

This section analyzes two single-sided mechanisms that focus on approximating the socially optimal allocation outcome. Although the approach of Zaman and Grosu (2013) respects quality awareness to some extent, the requirements of individual rationality and budget balance cannot be assessed. In contrast, the mechanism proposed by Karaenke (2014) guarantees individual rationality and budget balance but does not internalize quality as an intrinsic part of the mechanism.

Combinatorial auction-based allocation of virtual machine instances in Clouds

Zaman and Grosu (2013) study combinatorial auction mechanisms for allocating instances of virtual machines (VM) in Cloud computing. A VM abstracts the underlying physical computing resources from the customers by providing the VM as an electronic service of a specific QoS (e.g., storage size). The work of Zaman and Grosu (2013) is based on prior research in information systems (Archer et al., 2005; Lehmann et al., 2002). In their model, multiple customers demand instances of different types of VM offered by a single Cloud provider. To allocate VM, two distinguished mechanisms are proposed. The first mechanism, namely CA-LP, extends the work of Archer et al. (2005) by allowing customers to demand more than one VM instance of a given type. The winner determination problem is represented in the form of a linear program. The second mechanism, namely CA-GREEDY, is informed by the approach of Lehmann et al. (2002), which determines the allocation of VM instances based on the valuations of the customers and the total number of instances they demand. In the CA-GREEDY mechanism, the final allocation is approximated by a heuristic that implements a greedy algorithm. Zaman and Grosu (2013) advance this approach by integrating the relative sizes of the VM instances and show that the efficiency properties of the original mechanism continue to hold in the extended mechanism. With regard to the economic properties of their allocation outcome, they analyze both mechanisms for incentive compatibility. To verify the allocation efficiency achieved by these mechanisms, a simulation-based comparison with a real-world fixed-price mechanism is reported.

By analyzing the monotonicity properties of the proposed algorithms, Zaman and Grosu (2013) find that both CA-LP and CA-GREEDY induce customers to report their valuation for VM instances truthfully in expectation. In mechanism design theory, monotonicity properties are commonly used to establish incentive compatibility of a given mechanism. To this effect, the incentive compatibility constraints derived in this thesis critically depend on similar monotonicity properties (cf. Section 4.1.2

and Section 4.2.2). The model of Zaman and Grosu (2013) is instantiated in a simulation study by executing the artifact with artificial data. The associated experimental results indicate that both CA-LP and CA-GREEDY outperform the fixed-price mechanism in terms of resource utilization, generated revenue, and allocation efficiency. In particular, the CA-LP yields the highest revenue across all experiments. Because the linear program has to be solved repetitively, CA-LP is practically infeasible for a large number of customers. Instead, CA-GREEDY provides a reasonable alternative for Cloud computing settings with many customers.

The approach of Zaman and Grosu (2013) is subject to the following assumptions. Both allocation mechanisms CA-LP and CA-GREEDY are restricted to a single Cloud provider. In particular, whether the proven incentive compatibility of CA-LP and CA-GREEDY continues to hold when multiple Cloud providers enter the market remains an open question. From a mechanism design perspective, their model does not provide sufficient constructs to assess individual rationality or budget balance properties. In particular, their work does not provide results regarding the asymptotic efficiency of CA-GREEDY's approximated outcome relative to the outcome achieved by the optimal allocation.

Although the study of Zaman and Grosu (2013) is based on these assumptions, it is relevant to the mechanism in this thesis because Zaman and Grosu focus on allocating VM instances that have differentiated quality characteristics. These electronic services are valued differently by each provider. This thesis adopts this concept by defining the customers' valuations and the providers' costs as functions of their differentiated QoS demands and offers (cf. Section 3.2.2). Moreover, the CA-GREEDY mechanism by Zaman and Grosu determines the final allocation of VM instances by first ranking the customers by their valuation in decreasing order and then allocating them from top to bottom. In this thesis, a similar ranking technique can be used to determine the optimal matching between customer and providers by employing the concept of *positive assortative matching* (cf. Section 4.2.1). In order to determine the second-best mechanism that satisfies incentive compatibility, however, this assortative matching technique must be extended by imposing appropriate constraints that ensure incentive compatibility. Imposing such constraints results in a *truncated* positive assortative matching mechanism that entails truthful behavior of customers and providers (cf. Section 4.2.1).

Multiagent resource allocation in service networks

Karaenke (2014) studies the problem of allocating interdependent software services in agent-based multi-tier service networks. Multiple customer agents demand software services from multiple service provider agents that must reach agreements by negotiation. The proposed multi-tier resource allocation mechanism ensures that agreements are either established with the customer agents and with respective supplier agents in all tiers, or no agreements are established at all. The approach takes an interaction protocol engineering perspective from multiagent systems research (Huget and Koning, 2003) and uses models and methods from game theory (Parsons and Wooldridge, 2002) to study the mechanism properties. The resulting negotiation mechanism specifically avoids overcommitment by provider agents and is incentive compatible for a single customer agent demand. In addition, it constitutes a polynomial time heuristic for the social welfare maximization problem. Under the assumption that resources are non-substitutable between service providers, the mechanism arrives at a socially optimal allocation for single customer demands. The model of Karaenke (2014) is instantiated in a simulation study by executing the artifact with artificial data. He reports on a simulation study with and without substitutable resources where provider agents make decisions regarding the use of their own resources or the procurement of sub-services from provider agents in the next tier (in the case of substitutable resources). In particular, Karaenke (2014) studies the mechanism's efficiency for three distinct bidding policies: Providers either offer their internal resources first, external resources first or they offer best price resources only. The simulation study compares these three bidding policies by introducing the concept of the *utility ratio*. Utility ratios can be used to measure the overall efficiency of an approximated or second-best mechanism. This efficiency measure is defined as the ratio between the outcome of the second-best mechanism divided by the outcome of the optimal mechanism if it were to exist. Here, the outcome of the mechanism is defined as the sum of all utilities.

Using a game-theoretic framework, Karaenke (2014) finds that Nash equilibria exist for the optimal allocation of software services in multi-tier service networks. The underlying allocation mechanism is found to satisfy individual rationality and budget balance. For single customer demands, Karaenke (2014) finds that the proposed mechanism always arrives at the socially optimal allocation, given that resources are non-substitutable among providers. Once the number of customers increases and providers offer their external resources first, the efficiency of the mechanism is about 90% for substitutable resources. For non-substitutable resources, the efficiency of the mechanism decreases as the number of provider tiers increases. This finding suggests that the growing allocation complexity and solution space has a negative impact on the efficiency of the allocation mechanism. From a computational tractability perspective, it is found that computing socially optimal allocations is NP-complete.

The approach of Karaenke (2014) is based on the following assumptions. Although the proposed game-theoretic model facilitates the study of the economically desirable properties in service networks, it lacks the integration of two fundamental requirements that are necessary for the allocation of electronic services. First, the approach assumes single-sided competition on the market. In particular, the proposed mechanism does not satisfy incentive compatibility once additional customers enter the market. Second, Karaenke makes the assumption that merely homogeneous services are allocated by the mechanism. Thus, his model does not explicitly integrate the differentiated QoS desired by customers and offered by providers. The latter is in contrast to Zaman and Grosu (2013), who focus on electronic services of distinct QoS characteristics.

Although Karaenke's study is subject to these assumptions, it is relevant to this thesis because Karaenke uses the notion of *utility ratios* to measure the performance of the mechanism relative to the optimal allocation outcome. However, utility ratios are useful only if an optimal mechanism exists and is computationally feasible. Sandholm (2002) reports that determining the winners of combinatorial auction mechanisms is generally NP-complete. In this thesis, the winner determination problem is not NP-hard. Nonetheless, standard impossibility theorems from mechanism design prevent the existence of ex post optimal mechanisms that satisfy incentive compatibility and individual rationality (cf. Section 2.1.2.4). Therefore, unlike the mechanism of Karaenke, the mechanism proposed in this thesis is a second-best mechanism (cf. Section 3.6). Using the concept of utility ratios, the outcome of the second-best mechanism is compared to the outcome of the associated optimal mechanism if it were to exist. To obtain the outcome of the (non-existing) optimal mechanism in this thesis, the constraints on incentive compatibility and individual rationality are dropped. Then the efficiency of the mechanism is calculated as the ratio between the second-best and first-best outcome.

2.3.3 Double-sided, quality-unaware

This section focuses on two allocation mechanisms that support double-sided competition but do not internalize differentiated QoS. The first approach is concerned with mechanisms that face a tradeoff between incentive compatibility and optimality (Regev and Nisan, 2000). The second mechanism, presented by Gresik and Satterthwaite (1989), is a second-best mechanism that satisfies Bayesian incentive compatibility, individual rationality, and budget balance.

The POPCORN market – an online market for computational resources

Regev and Nisan (2000) study a set of market-based mechanisms for trading central processing unit (CPU) cycles across the Internet. CPU cycles of idle processors are offered as electronic services, whose QoS is specified by the number of available CPU cycles. In the electronic market studied by Regev and Nisan (2000), multiple customers demand CPU cycles, while multiple providers offer CPU cycles at a specific price. Their model is based on prior research in market-based resource allocation (Waldspurger et al., 1992). In their work, the market itself is responsible for matching customers and providers according to economic criteria. The two main design goals of the system are to implement economically efficient allocations by maximizing the global utility and to motivate customers and providers to reveal their true utility of the CPU cycles. The model of Regev and Nisan is instantiated in a simulation study by executing the artifact with artificial data. In this study, they focus on three distinct sealed-bid mechanisms. First, they consider a repeated Vickrey auction where the price paid by the customer corresponds to the second highest price offered in the auction (Vickrey, 1961). Second, they study a simple double auction where customers and providers offer lower and upper bounds of a price. This auction is similar to the simultaneous execution of two first-price sealed bid auctions (Milgrom, 2004). Each provider places an offer at a high price, which is then automatically decreased at a given rate until a customer is found. Similarly, each customer submits a low price, which is then automatically increased at a given rate until a provider is found. Third, Regev and Nisan study a repeated clearinghouse double auction that calculates the demand and supply curves and sets the equilibrium price at fixed time intervals. Every transaction in the current round uses the same equilibrium price. For each mechanism implementation, Regev and Nisan study the asymptotic efficiency and the price stability of the associated allocation outcome. Furthermore, their model has been implemented and was operable for over one year through a website (POPCORN project, 1998).

Regev and Nisan (2000) find that the clearinghouse mechanism achieves an efficiency of about 96%. The repeated Vickrey auction and the simple double auction decrease in efficiency, falling below 85% and 70%, respectively. This finding shows the effect of the concurrency of these two auction formats on the allocation efficiency. Low-cost providers who arrive later cannot be matched to high-valued customers anymore because they have left the market already. The opposite applies to expensive providers. In terms of price stability, Regev and Nisan find that transaction prices generated by the clearinghouse auction are more stable than those from other auction formats. Regarding the probability of trade, their findings confirm that high offers of customers and low prices of providers lead to an increased probability of trade in all three auction formats.

The study of Regev and Nisan makes the following assumptions on the proposed mechanism implementations. First, Regev and Nisan do not explicitly internalize differentiated QoS in the mechanism. Thus, their model cannot be readily used to analyze mechanisms in which traders demand and offer services of distinct QoS. Second, although the single-round Vickrey auction is incentive compatible (even in dominant strategies), incentive compatibility does not necessarily hold if it is executed repeatedly. Furthermore, the repeated Vickrey auction does not support competition among the providers. In contrast, the simple double auction supports (Bayesian) incentive compatibility among both customers and providers. However, it does not necessarily arrive at the socially optimal allocation. Lastly, because Regev and Nisan assume multi-unit resources, the clearinghouse double auction is not incentive compatible.

Although the three allocation mechanisms studied by Regev and Nisan (2000) are subject to these assumptions, they are relevant to the approach pursued in this thesis because Regev and Nisan's clearinghouse double auction also uses a single-round matchmaking mechanism to determine the optimal allocation. In order to obtain a mechanism that integrates differentiated QoS on both market sides, however, this matchmaking mechanism must internalize these differentiated QoS characteristics. Therefore, this thesis drops the assumption of homogeneous QoS by defining an allocation rule that matches customers and providers based on the reported bids about their differentiated QoS demands and offers. This modification ensures that the mechanism matches the right pairs of customers and providers, thus eliciting the highest possible match surplus (cf. Section 4.2.1).

The rate at which a simple market converges to efficiency as the number of traders increases: an asymptotic result for optimal trading mechanisms

Gresik and Satterthwaite (1989) study the replication of the bilateral model proposed by Myerson and Satterthwaite (1983) to allow for multiple customers and providers with double-sided private information. Customers and providers exchange identical objects based on privately known reservation values. Using a mechanism design framework, Gresik and Satterthwaite characterize the set of all allocation mechanisms that satisfy incentive compatibility and individual rationality. Because ex post optimality is not attainable by the Myerson-Satterthwaite impossibility theorem, Gresik and Satterthwaite construct an ex ante optimal mechanism. In their work, an ex ante efficient mechanism is one that places equal welfare weights on every trader and maximizes the sum of the traders' ex ante expected utilities subject to incentive compatibility and individual rationality. Gresik and Satterthwaite provide conditions sufficient for the existence of ex ante efficient mechanisms and study specific convergence properties of these mechanisms as the market size increases. In particular, they determine the rate at which the ex ante optimal mechanism converges to ex post optimality as a function of the market size. The model of Gresik and Satterthwaite is instantiated in a numerical example, which illustrates the convergence properties of the ex ante optimal mechanism for an increasing market size up to twelve traders. Finally, Gresik and Satterthwaite compare the efficiency properties of the optimal mechanism to those of a fixed-price mechanism.

Gresik and Satterthwaite (1989) find that their characterization of all incentive compatible and individually rational mechanisms in the multilateral model generalizes the results obtained in the bilateral trading model of Myerson and Satterthwaite (1983). The numerical example for uniformly distributed reservation values shows that the bilateral model achieves an efficiency of 84.36%. These inefficiencies disappear with increasing market size. In the limit market, as the number of customers and providers approaches infinity, the mechanism converges to perfect competition; that is, ex post efficiency is achieved. Gresik and Satterthwaite also provide an upper bound on the relative inefficiency of the mechanism. For the fixed-price mechanism, Gresik and Satterthwaite find that the associated convergence rate is significantly lower than that obtained for the optimal trading mechanism.

The study of Gresik and Satterthwaite (1989) is subject to the following assumptions. Gresik and Satterthwaite study the allocation of identical objects only; that is, all providers offer identical objects and all customers desire identical objects. Hence, their model cannot address the allocation of differentiated objects such as electronic services with distinct QoS requirements. If only homogeneous objects are considered in the model, the optimal match of customers and providers cannot be determined because it does not matter which customer trades with which provider. As a consequence, the model presented by Gresik and Satterthwaite assumes reservation values that only depend on the trader's private information. In such a case, however, the mechanism cannot elicit the highest possible match surplus.

Although the mechanism developed by Gresik and Satterthwaite (1989) assumes the allocation of identical objects, it is relevant to the mechanism presented in this thesis because Gresik and Satterthwaite characterize a direct revelation mechanism for allocating objects with double-sided private information. Subsequent approaches adopt this method to derive second-best mechanisms whose efficiency properties have been studied successfully (Zacharias and Williams, 2001; Cripps and Swinkels, 2006). This thesis also invokes the revelation principle to determine the optimal allocation rules (cf. Section 3.3). Similar to Gresik and Satterthwaite, the associated expected payments guarantee Bayesian incentive compatibility among all traders (cf. Section 4.2.1). In order to support the allocation of services with differentiated QoS, this thesis must extend Gresik and Satterthwaite's utility functions by internalizing the the QoS desired by customers and the QoS offered by providers (cf. Section 3.2.2). Consequently, this extension supports the allocation of electronic services with distinct QoS requirements on both market sides.

2.3.4 Double-sided, incentive incompatible

This section provides an analysis of three double-sided mechanisms that respect quality awareness but are not incentive compatible. While the approach by Denoeud-Belgacem et al. (2010) does not consider incentive compatibility whatsoever, the work of Schnizler et al. (2008) estimates efficiency losses that emerge due to incentive incompatibility. Samimi et al. (2016) extend the single-sided mechanism by Zaman and Grosu (2013) to double-sided, quality-aware mechanisms.

Combinatorial auctions for exchanging resources over a Grid network

Denoeud-Belgacem et al. (2010) study combinatorial auctions for exchanging resources over a Grid network. Their model is based on prior research in management science (Kwasnica et al., 2005). A Grid network is an interconnection between computers, which provides access to electronic services in the form of computing resources such as processing power, storage space or applications. On the marketplace studied by Denoeud-Belgacem et al. (2010), multiple customers and providers submit bids for trading computational resources. They consider two types of resources, namely computing and storage resources, whose quantity and quality are specified. Customers and providers desire/offer combinations of these resources. Providers submit their bids as bundles of resources available during a given time frame. These bundles contain the reservation price the providers are willing to accept. Customers express their bids by specifying their desired quality attributes and the maximum price they are willing to pay. The proposed winner determination algorithm matches the desired/offered resource requirements and determines the allocation that maximizes the social welfare among all valid matches. Because the combinatorial winner determination problem is NP-hard, Denoeud-Belgacem et al. (2010) propose an iterated greedy heuristic that allocates resources in a “first come, first serve” manner. With regard to pricing, Denoeud-Belgacem et al. (2010) study a variant of the market clearing price approach commonly used in double auctions (Wilson, 1985). In particular, they approximate the market clearing price by minimizing the gaps between the losing bids and the computed prices of the bundles in case the market clearing property is not verified. This pricing mechanism guarantees individual rationality and budget balance. Finally, the model of Denoeud-Belgacem et al. (2010) is instantiated in a simulation study by executing the artifact with artificial data. In this study, they compare the mechanism’s performance of the exact solution with the approximated solution of the winner determination problem. They also provide experimental results on the different pricing schemes.

Denoeud-Belgacem et al. (2010) find that due to the exponential complexity of the winner determination problem, computing the optimal solution is infeasible. In a setting with 100 customers and 100 providers in different time slots, the proposed greedy heuristic arrives at an asymptotic efficiency of about 86% compared to the optimal solution. Denoeud-Belgacem et al. find that the proposed pricing schemes tend to favor the providers because they receive a higher share of the generated surplus.

The approach of Denoeud-Belgacem et al. (2010) is based on the following assumptions. First, their mechanism assumes a pricing model that does not support incentive compatibility. Another drawback is their incomplete formal framework, which is necessary for a theoretically sound assessment of the mechanism design properties such as budget balance and individual rationality. As a result, it remains unclear to what extent the proposed greedy algorithm satisfies these properties. Furthermore, the asymptotic efficiency of their mechanism for large markets cannot be properly assessed.

Despite these limitations, Denoeud-Belgacem et al. list a set of economic requirements on the allocation mechanism that are relevant to the approach in this thesis. Due to the standard impossibility theorems from mechanism design theory, some of these requirements cannot be fulfilled simultaneously. Therefore, Denoeud-Belgacem et al. propose a tradeoff between these requirements: Their greedy mechanism gives up incentive compatibility. Abandoning incentive compatibility, however, is critical because it encourages the traders to strategically misrepresent their true reservation values. Therefore, this thesis drops the assumption of incentive incompatible payments by imposing appropriate constraints for incentive compatibility on the winner determination problem.

Trading Grid services – a multi-attribute combinatorial approach

Schnizler et al. (2008) study a multidimensional combinatorial auction mechanism for trading electronic services among multiple providers and customers in a Grid network. Their model is informed by prior research in auction theory (Milgrom, 2004) and Grid technologies (Lai et al., 2005). The traded services respect resource functionalities (e.g., storage) and quality characteristics (e.g., size), as well as dependencies and time attributes. For the design of the allocation mechanism, Schnizler et al. focus on common desirable properties from mechanism design theory including incentive compatibility, individual rationality, and budget balance. In particular, the mechanism supports double-sided competition as a domain-specific requirement for Grid environments, where multiple service providers publish their services and multiple service customers discover them. The bidding language of the mechanism allows customers and providers to specify their desired and offered QoS and includes combinatorial bids on these attributes. Based on this multi-attribute bidding language, Schnizler et al. represent the combinatorial winner determination problem as a linear mixed integer problem in order to achieve a socially optimal solution. Subsequently, the authors study two different pricing schemes to determine the payments for the allocated Grid services. First, they consider payments based on the prominent Vickrey-Clarke-Groves (VCG) mechanism (Vickrey, 1961). Using these payments, the proposed mechanism is incentive compatible and individually rational but runs a permanent deficit in funds. Second, they study the k-pricing scheme, in which prices are calculated based on the difference between the bids of customers and providers (Satterthwaite and Williams, 1993). Furthermore, the model of Schnizler et al. is instantiated in a simulation study by executing the artifact with artificial data. By means of this experimental evaluation, Schnizler et al. study the effect of the k-pricing scheme on the strategic behavior of customers and providers.

Schnizler et al. (2008) find that the k-pricing scheme does not sufficiently punish strategic misrepresentation of the traders' bids. This finding is a result of the Myerson-Satterthwaite impossibility theorem discussed in Section 2.1.2.4, which plays a crucial role in this thesis as well. Using the VCG outcome as a benchmark and given the set of probability distributions, the mechanism based on the k-pricing scheme achieves an efficiency of approximately 41.52%, averaged over 10 or 20 agents. The simulation study also indicates that the traders' misreporting is limited to small markets. As the number of misreporting traders increases, so too does the risk of not being allocated by the mechanism. Regarding computational tractability, Schnizler et al. find that the proposed mechanism is computationally very demanding. Approximate solutions, however, entail adequate runtime results for up to 500 traders.

The study of Schnizler et al. (2008) makes the following assumptions regarding the design and the evaluation of their allocation mechanism. Similar to the work of Denoeud-Belgacem et al. (2010), Schnizler et al. are faced with a tradeoff between the four desirable economic properties of incentive compatibility, individual rationality, budget balance, and ex post optimality. Because they assume a k-pricing scheme, the mechanism fails to implement incentive compatibility. Hence, the allocation outcome of the mechanism is inefficient. Although the proposed simulation study provides a set of results on the mechanism's efficiency losses, Schnizler et al. do not provide details regarding the asymptotic efficiency of the mechanism for limit markets. Thus, it remains unclear whether the mechanism's inefficiencies are bounded when the number of agents becomes large.

Despite these assumptions, the model of Schnizler et al. (2008) is of special interest for this work because Schnizler et al. consider a double-sided market for trading electronic services of specific QoS requirements. In particular, Schnizler et al. study the efficiency properties by using the mechanism's optimal outcome as a benchmark. In their work, the optimal outcome arises from a VCG implementation of the mechanism. It is well known in mechanism design theory that the VCG auction is a direct revelation mechanism (Fudenberg and Tirole, 1993). Subsequent approaches for electronic service allocation in double-sided markets use the outcome of direct revelation mechanisms as a benchmark to quantify the efficiency of approximated mechanisms and heuristics (Stösser et al., 2010). This thesis adopts this approach to measure the inefficiencies caused by the asymmetry of information about the agents' QoS (cf. Section 3.2.4). For designing a second-best mechanism that satisfies incentive compatibility, however, the mechanism presented in this thesis uses transfer functions that emerge from invoking the revelation principle. In contrast to Schnizler et al.'s k-pricing scheme, these transfers satisfy incentive compatibility (cf. Section 4.1.2).

A combinatorial double auction resource allocation model in Cloud computing

Samimi et al. (2016) study a combinatorial double auction for Cloud computing based on the approach for single-sided allocation mechanisms described by Zaman and Grosu (2013). The extended model allows both providers and customers to compete for resources on the basis of QoS demands. The proposed combinatorial double auction considers the Cloud service requirements of both sides of the market and permits bidding for resource combinations. Compared to allocation mechanisms based on the single-item auction, their model performs more efficiently because the agents are able to bid for multiple items. Samimi et al. describe their allocation mechanism in the following sequence. First, customers and providers submit their bids to the mechanism. The auctioneer generates bundles consisting of the required resource combinations and closes the auction. Next, the auctioneer determines the final allocation of resources. The underlying winner determination algorithm is based on a kind of assortative matching, where the bids are ranked against each other in descending order. For the matched customer-provider pairs, the auctioneer applies a pricing model based on the *k*-pricing scheme (Satterthwaite and Williams, 1993) to compensate providers for Cloud service delivery. The model of Samimi et al. is instantiated in a simulation study by executing the artifact with artificial data. They also report on the economic efficiency and incentive compatibility properties of their allocation mechanism.

Samimi et al. (2016) find that their allocation mechanism yields an economically efficient outcome. Their notion of economic efficiency refers to the magnitude of utility that customers and providers can receive. They find that providers with low bids and customers with high bids obtain the highest utilities. In contrast, providers with high bids and customers with low bids even obtain negative utilities. Samimi et al. also find that customers and providers are likely to report their bids truthfully. Hence, they claim that their mechanism satisfies incentive compatibility.

The approach of Samimi et al. (2016) has the following shortcomings. First, the assumptions made in the model are not sufficiently disclosed. Samimi et al. present their findings by interpreting a set of numerical results that are difficult to replicate and verify. Mathematical proofs are not used to support their statements. Second, the notion of economic efficiency is not sufficiently explained. In particular, Samimi et al. claim that their mechanism induces truthful behavior, yet formal proofs for incentive compatibility are omitted. In fact, *k*-pricing mechanisms are generally not incentive compatible.

Although the study of Samimi et al. (2016) is based on these assumptions, it is interesting that Samimi et al. also use a ranking technique of the submitted bids to determine the final allocation of electronic services. As discussed earlier, Zaman and Grosu (2013) used a similar ranking technique, which is, however, limited to a single provider. Because double-sided competition is supported by Samimi et al., the mechanism sorts the bids of customers in descending order and the bids of providers in ascending order. Then, the mechanism matches customers and providers that are on the same rank from top to bottom. This procedure is a version of a *negative assortative* allocation scheme, in which the highest value on one side is matched to the lowest value on the other side of the market (Shimer and Smith, 2000). For determining the optimal allocation of services with distinct QoS, however, the assortative allocation rule must be redefined and enriched with appropriate incentive compatibility constraints. Therefore, the mechanism proposed in this thesis uses a positive assortative matching rule that imposes additional constraints to ensure incentive compatibility and individual rationality (cf. Section 4.2.1).

2.3.5 Double-sided, incentive compatible

This section analyzes two double-sided allocation mechanisms that satisfy incentive compatibility and support quality awareness. While the mechanisms proposed by Muthoo and Mutuswami (2005, 2011) do not derive the socially optimal allocation rules and payments, the matching mechanism in Johnson (2013) considers an optimality criterion from the perspective of a profit-maximizing intermediary.

Competition and efficiency in markets with quality uncertainty / Imperfect competition and efficiency in lemons markets

The integration of uncertainty about an object's quality characteristics into the design of allocation mechanisms was originally proposed by Akerlof (1970). On the so-called lemons market, customers and providers exchange lemons of privately known qualities. Muthoo and Mutuswami (2011) study the efficiency properties of an allocation mechanism with private information about the quality of the traded objects. Their model is based on prior research in economic theory (Akerlof, 1970; Samuelson, 1984). Each provider has private information about the quality of a single, indivisible object. On the other hand, there is a single customer who demands one of these objects. Providers offer their objects at two distinct quality levels, namely high or low. Similar to Gresik and Satterthwaite (1989), Muthoo and Mutuswami (2011) invoke the revelation principle to derive a direct revelation mechanism which satisfies incentive compatibility and individual rationality. Because private information about quality is present in their analysis, this mechanism is a second-best mechanism. In particular, Muthoo and Mutuswami (2011) study the efficiency properties of this mechanism for a varying number of competing providers. In an earlier study, Muthoo and Mutuswami extended the model by allowing for multiple customers who have private information on their distinct quality needs. Using a numerical analysis, they studied the efficiency properties of the second-best mechanisms that emerge for such markets with double-sided private information about quality.

Muthoo and Mutuswami (2011) find that increasing competition among providers may not cause the mechanism to achieve full efficiency in some situations. In particular, if the number of providers remains below a specific threshold, the efficiency of the second-best mechanism stagnates. Thus, the optimal number of providers in such markets is finite. For the extended model with double-sided information (multiple providers and multiple customers), Muthoo and Mutuswami (2005) find that the second-best mechanism achieves full efficiency if the probability that the customer desires low quality is greater than 55%. If this probability is 50%, however, the mechanism is inefficient even when the number of providers becomes arbitrarily large. Unexpectedly, the efficiency of the mechanism appears to be non-monotonic if the probability that the customer desires low quality is 40%. In addition, their findings indicate that allowing for private information on the customers' side can be used to relax the incentive compatibility constraint of low quality providers.

The studies of Muthoo and Mutuswami (2005, 2011) are subject to the following assumptions regarding the design of their mechanisms. Although Muthoo and Mutuswami (2011) derive a second-best mechanism with double-sided private information about quality to study the mechanism's efficiency properties, their model assumes a single customer. The extended model proposed in Muthoo and Mutuswami (2005) explicitly allows for multiple customers with private quality information. Yet in this extension, Muthoo and Mutuswami (2005) do not derive the optimal allocation rules or payments used by the second-best mechanism. Hence, it remains unclear how the designer must determine the transfers so as to incentivize truthful behavior. Moreover, private information about the quality of the objects is limited to two discrete quality levels (i.e., high and low). Therefore, the model of Muthoo and Mutuswami (2005, 2011) does not support arbitrary (i.e., continuous) quality values.

Despite these assumptions, their studies are relevant to this thesis because Muthoo and Mutuswami explicitly allow for differentiated quality characteristics of the traded objects. In particular, these quality characteristics are modeled as the private information of each trader in the market. Subsequent research applies this concept by integrating the object's quality as private information for the special case of dynamic matching markets (Kultti et al., 2015). This thesis adopts this approach by modeling each trader's distinct quality requirements as the private information of that trader (cf. Section 3.2). In order to obtain a mechanism that allocates services of differentiated QoS, however, this thesis drops the assumption of discrete quality levels by defining the QoS as continuous private information. Moreover, Muthoo and Mutuswami provide the boundaries for allocation in the first-best and the second-best mechanism. These allocation boundaries illustrate the distortion due to the presence of private information about quality. Because this thesis extends the work of Muthoo and Mutuswami toward continuous QoS, the associated allocation boundaries must be adapted as well (cf. Section 4.2.3).

Matching through position auctions

If the social welfare depends on the privately known QoS of the matched customers and providers, double-sided *matching mechanisms* may be used to determine the allocation of electronic services. In these mechanisms, customers and providers only produce mutual surplus if they are matched together. With a setting of double-sided private information, Johnson (2013) studies the optimal matching and payment rules from the perspective of a profit-maximizing intermediary. The direct revelation mechanism employs the concept of positive assortative matching; that is, agents of the highest quality level are matched together, the second highest likewise, and so on. This mechanism is incentive compatible and individually rational for all agents. One way to implement this mechanism is by position auctions, in which multiple customers and providers simultaneously submit their bids to the mechanism, the bids are ranked against each other and the agents are allocated accordingly. Johnson studies two different formats used in position auctions, namely the winners-pay format and the all-pay format. In the winners-pay format only matched agents pay their bid to the auctioneer. The all-pay format requires all agents to pay their bid, regardless of a successful match. Finally, Johnson analyzes the monotonicity properties of the bid functions and the effect of the market size on the implementability of the profit-maximizing allocation in the winners-pay format.

Johnson (2013) finds that his model with multiple customers and providers generalizes some of the results obtained by Myerson and Satterthwaite (1983) on how an intermediary engages in price discrimination. In addition to identifying standard conditions sufficient for maximizing the intermediary's profit, Johnson finds that a set of hazard rate bounds is necessary to guarantee that higher types are eligible to match to larger sets of partner qualities. With these assumptions, it suffices to solicit a one-dimensional bid from the agents to determine the profit-maximizing match and to implement the profit-maximizing payments. In the case of the winners-pay format, however, Johnson finds that the bid functions must be increasing to enable the implementability of the mechanism. Finally, it is found that in sufficiently large markets, the non-monotonicities in the bid functions disappear.

The study of Johnson (2013) is based on the following assumptions. Johnson assumes that the matching mechanism is used by a profit-maximizing intermediary. Thus, he does not derive the optimal allocation mechanism and payments from the perspective of a social planner, who seeks to maximize the expected social welfare. Because the intermediary's profit is maximized by the mechanism, it does not facilitate distributed decision-making among the agents. Moreover, Johnson does not study the asymptotic efficiency properties of the optimal allocation rules.

The approach taken in this thesis is informed by the work of Johnson (2013). More specifically, the allocation mechanism proposed in this thesis extends Johnson's mechanism by maximizing the expected social welfare under budget balance instead of the intermediary's profit. This change requires re-defining the optimization problem faced by the mechanism (cf. Section 4.2). In Johnson's model, the agents on both market sides are of the same "type" in the sense that they all use the same valuation function. In contrast to this model, markets for allocating electronic services involve two distinct types of agents, namely customers and providers. Consequently, the model proposed in this thesis extends that of Johnson by supporting customer-provider matching in double-sided markets. In this extension, the agents' valuation function defined by Johnson remain unchanged for customers. For providers, however, a cost function must be defined that integrates the differentiated QoS of the matched agent pairs (cf. Section 3.2.2). Once these modifications are in place, the approaches of Gresik and Satterthwaite (1989) as well as Muthoo and Mutuswami (2005, 2011) can be applied to study the asymptotic efficiency properties of the emerging second-best allocation mechanism.

Chapter 3

Model

For He said, "See that you make all things according to the pattern shown you on the mountain."

Hebrews 8:5b, The Bible
(New King James Version)

This chapter describes the formal model used for designing the allocation mechanism for electronic services with private quality information. The framework is based on the notation introduced in Chapter 2 as well as on the model of Johnson (2013), who derived a set of conditions for profit-maximizing matching markets. This section represents each requirement identified in Section 2.2 in the formal model.

3.1 Representing double-sided competition

The allocation mechanism for electronic services must support double-sided competition among multiple customers and multiple providers. Thus, the model must allow both customers and providers to submit their QoS demands and offers to the mechanism. Let $\mathscr{A}_C$ with $N = |\mathscr{A}_C|$ denote the set of all customers and let $\mathscr{A}_P$ with

$M = |\mathscr{A}_P|$ denote the set of all providers. Each customer $a_i \in \mathscr{A}_C$ demands a single, standardized electronic service by submitting a bid that specifies his desired QoS. For instance, customers demand standardized storage services by submitting the amount of the required storage. Likewise, each provider $b_j \in \mathscr{A}_P$ offers a single electronic service by submitting a bid that specifies the QoS he actually offers. As discussed in Section 2.2.2, standardized services are electronic services of equal functional properties. Since M competing customers demand services from N competing providers, the maximum number of allocations that can arise in this setting is given by $K = \min\{M, N\}$. To extend bilateral models for service allocation, it is assumed that $K \geq 2$.

3.2 Representing quality as private information

In this section, the QoS desired by customers and the QoS offered by providers are internalized in the allocation mechanism. Because the quality of the services are mutually unknown, it is modeled as double-sided private information. This private information reflects the preferences each agent assigns to the electronic service. To this end, valuation and cost represent the agents' reservation values for the electronic services. Since customers and providers only produce mutual surplus if they are matched together, valuation and cost functions depend on the privately known QoS of each agent. Using these interdependent functions, the expected utilities of the agents can be defined. Finally, this section presents the definitions of the social choice function and the direct revelation mechanism of the allocation problem in this research.

3.2.1 Desired and offered quality

Each customer privately observes the desired quality θ_i that he has assigned to the demanded service. Because desired quality is that customer's private information, it is drawn from a probability density function $f(\theta_i)$, being strictly positive on $[0, \bar{\theta}]$

with cumulative distribution function $F(\theta_i)$. On the supply side, each provider privately learns the actual quality σ_j that he can offer for the electronic service. The quality offered by providers is drawn from a probability density function $h(\sigma_j)$, being strictly positive on $[0, \bar{\sigma}]$ with cumulative distribution function $H(\sigma_j)$. Let $\theta = (\theta_1, \ldots, \theta_N)$ and $\sigma = (\sigma_1, \ldots, \sigma_M)$ be the private information vectors of all customers and providers, respectively. The vectors $\theta_{-i} = (\theta_1, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots, \theta_N)$ and $\sigma_{-j} = (\sigma_1, \ldots, \sigma_{j-1}, \sigma_{j+1}, \ldots, \sigma_M)$ specify the private information except for a_i and b_j, respectively. Further, let $\mathbb{E}_{\theta_{-i},\sigma}$ denote the expectation over all private information conditional on customer a_i's information, and let $\mathbb{E}_{\theta,\sigma_{-j}}$ denote the analogous expectation for provider b_j. The unconditional expectation over all private information is denoted by $\mathbb{E}$.

3.2.2 Valuation and cost

Customer a_i's valuation for consuming a service is given by $v(\theta_i, \sigma_j)$. This valuation can be interpreted as a_i's maximum willingness to pay for the service. It depends on a_i's desired quality θ_i, as well as on the difference between his own desired quality θ_i and provider b_j's offered quality σ_j. The definition of $v(\theta_i, \sigma_j)$ captures the fact that provider b_j does not know a_i's willingness to pay. This research focuses on customer valuation functions that are non-monotonic in the quality offered by the provider. For instance, a customer may prefer a service with medium over high computational capacity. A high-capacity service is well able to process many simultaneous requests from the customer's application. If, however, this application is not provided with enough computational power or resources, it will fail to answer these simultaneous requests in due time. This leads to higher buffering in the application and thus longer response times. Thus, a customer's valuation must take into account the application that uses the service and the tradeoff between being idle or buffering heavily (Hang and Singh, 2010). Therefore, any mismatch between desired quality and offered quality creates adjustment problems for the customer. That is, $v(\theta_i, \sigma_j)$ is maximized when the supplied quality and the desired quality are equal (i.e., when $\theta_i = \sigma_j$). By assumption, the maximal value is increasing in θ_i.

On the supply side, each provider b_j's provision cost $c(\theta_i, \sigma_j)$ for his service depends on his offered quality σ_j and on the difference between his own offered quality σ_j and the customer's desired quality θ_i. This definition suggests that customers do not know the provision costs of providers. If a provider produces a quality lower than the quality desired by a customer, this provider incurs higher cost from not fulfilling the requirements. If, in contrast, a provider maintains higher quality than desired, his cost increases due to idle resources (Greenberg et al., 2009). Hence, $c(\theta_i, \sigma_j)$ is minimized when $\sigma_j = \theta_i$ and when the minimal value is increasing in σ_j. This assumption captures the fact that a mismatch in offered quality and desired quality creates higher provision costs resulting from after-sales customer service cost and missed opportunity cost. Both $v(\theta_i, \sigma_j)$ and $c(\theta_i, \sigma_j)$ are assumed to be thrice differentiable.

3.2.3 Expected utilities

As discussed in Section 2.1.1.1, this research assumes risk-neutral agents with quasi-linear utility functions. This utility function is the internalization of each agent's performance measure and is based on the preferences concerning the allocation of electronic services. Hence, utility functions are additively separable and linear in money and in the reservation value for the electronic service.

For describing an allocation of a service, a decision variable is used. Decision variables suit well to model the allocation of services between customers and providers. Let $x_{ij}(\theta, \sigma) \in [0,1]$ denote the probability that a service is allocated from provider b_j to customer a_i. For example, if $x_{11}(\theta, \sigma) = 1$, the service is provided by provider b_1 and consumed by customer a_1 in the final allocation. For ease of exposition, this model adopts vectorized function arguments without further reference and, for example, uses the convention $x_{ij}(\theta_i, \theta_{-i}, \sigma) \equiv x_{ij}(\theta_1, \ldots, \theta_N, \sigma_1, \ldots, \sigma_M)$ in any appropriate context.

Using this notation, each customer's expected utility function is given by

$$U_C(\theta_i) = \mathbb{E}_{\theta_{-i},\sigma}\Big[\sum_{j=1}^{M} v(\theta_i, \sigma_j) x_{ij}(\theta_i, \theta_{-i}, \sigma) - t_C(\theta_i, \theta_{-i}, \sigma)\Big], \tag{3.1}$$

where $t_C(\theta_i, \theta_{-i}, \sigma)$ is the monetary transfer made by a_i conditional on all other agents' private information. Similarly, the expected utility of provider b_j is given by

$$U_P(\sigma_j) = \mathbb{E}_{\theta,\sigma_{-j}}\Big[t_P(\theta, \sigma_j, \sigma_{-j}) - \sum_{i=1}^{N} c(\theta_i, \sigma_j) x_{ij}(\theta, \sigma_j, \sigma_{-j})\Big], \tag{3.2}$$

where $t_P(\theta, \sigma_j, \sigma_{-j})$ corresponds to the monetary compensation provider b_j receives for providing a service to customer a_i.

3.2.4 Mechanism definition

As discussed in Section 2.1.1, this research characterizes the allocation mechanism for electronic services in terms of a *direct revelation mechanism*. In a direct revelation mechanism, all customers are induced to reveal their desired QoS θ_i and all providers are induced to reveal their offered QoS σ_j to the mechanism. Then, the mechanism dictates the allocation of electronic services and the respective monetary transfers. The direct revelation mechanism is thus defined by the decision variables $x_{ij}(\theta, \sigma)$, as well as the monetary transfers that accrue for providing and consuming electronic services. As discussed in Section 2.1.1.2, the social choice function for this setting is given by

$$f_{sc} : \mathbb{R}^+ \times \cdots \times \mathbb{R}^+ \rightarrow \Omega \tag{3.3}$$

$$(\theta, \sigma) \mapsto (x_{ij}(\theta, \sigma), t_C(\theta, \sigma), t_P(\theta, \sigma)) \tag{3.4}$$

for all customers a_i and providers b_j. Using the definition for social choice functions, the direct revelation mechanism is given by a set of functions

$$\mathscr{M} = \{x_{ij}(\theta, \sigma), t_C(\theta, \sigma), t_P(\theta, \sigma)\} \tag{3.5}$$

for all a_i and b_j.

3.3 Representing incentive compatibility

To determine the optimal equilibrium outcome from among all possible mechanisms, this research adopts the standard approach in mechanism design theory by invoking the revelation principle (cf. Section 2.1.1.4). The revelation principle essentially states that for any outcome of a mechanism that is in Bayesian Nash equilibrium, there exists a direct revelation mechanism with an identical equilibrium outcome. As pointed out in Section 2.1.1.3, the set of equilibrium outcomes induced by direct revelation mechanisms is that which satisfies *Bayesian incentive compatibility*; that is, all agents' honest type reporting is a Bayesian Nash equilibrium of the underlying game. The concept of *Bayesian Nash equilibria* is applicable to this model because each agent's QoS is known only to that agent and is observed neither by the mechanism designer nor by the other agents.

Recall from Section 2.1.2.1, the mechanism is (Bayesian) incentive compatible if and only if for each customer's reported desired quality $\hat{\theta} \neq \theta_i$

$$U_C(\theta_i) \geq \mathbb{E}_{\theta_{-i},\sigma}\Big[\sum_{j=1}^{M} v(\theta_i,\sigma_j)x_{ij}(\hat{\theta},\theta_{-i},\sigma) - t_C(\hat{\theta},\theta_{-i},\sigma)\Big], \tag{3.6}$$

and for each provider's reported offered QoS $\hat{\sigma} \neq \sigma_j$

$$U_P(\sigma_j) \geq \mathbb{E}_{\theta,\sigma_{-j}}\Big[t_P(\theta,\hat{\sigma},\sigma_{-j}) - \sum_{i=1}^{N} c(\theta_i,\sigma_j)x_{ij}(\theta,\hat{\sigma},\sigma_{-j})\Big]. \tag{3.7}$$

3.4 Representing individual rationality

In addition to incentive compatibility, the mechanism must satisfy individual rationality. As discussed in Section 2.1.2.2, rational customers and providers will not participate in the mechanism if their expected utility is negative (evaluated at one of the distinct timing stages). This constraint is also known as the voluntary participation constraint, which ensures that no agent is worse off after the mechanism is run. Similar to common mechanism design settings, this research considers the *interim* timing state for individual rationality. At the interim stage, each agent can withdraw from the mechanism once he observes his own type, not knowing the types of all other agents. Thus, the individual rationality constraint ensures that, conditional on his private information, every agent imputes an expected utility from participating in the mechanism that is greater than or equal to the utility of his outside option. In this research, the outside option of each agent is zero. Formally, the mechanism is individually rational if and only if

$$U_C(\theta_i) \geq 0 \qquad \text{for customer } a_i \text{, and} \tag{3.8}$$

$$U_P(\sigma_j) \geq 0 \qquad \text{for provider } b_j. \tag{3.9}$$

3.5 Representing budget balance

As pointed out in Section 2.1.2.3, the mechanism must raise sufficient funds from the agents to cover the cost for running the platform. In this research, the mechanism balances the budget if all expected transfers made among the agents add to zero. This constraint is necessary because the mechanism must not depend on any external source of funds. That is, (ex ante) budget balance in this model is defined as the difference between the sum of all expected monetary compensations received by providers and the sum of all expected monetary payments made by customers for service allocation. Formally, (ex ante) budget balance is defined as

$$\begin{aligned}\sum_{i=1}^{N}\int_{0}^{\bar{\theta}}&\mathbb{E}_{\theta_{-i},\sigma}[t_C(\theta_i,\theta_{-i},\sigma)]f(\theta_i)d\theta_i\\&-\sum_{j=1}^{M}\int_{0}^{\bar{\sigma}}\mathbb{E}_{\theta,\sigma_{-j}}[t_P(\theta,\sigma_j,\sigma_{-j})]h(\sigma_j)d\sigma_j=0.\end{aligned}\tag{3.10}$$

3.6 Representing optimality

Since ex post optimality is unattainable when incentive compatibility, individual rationality, and budget balance are required, this research attempts to find a mechanism that achieves ex ante optimality. As discussed in Section 2.1.2.4, an ex ante optimal mechanism is a second-best mechanism that maximizes the expected social welfare among all incentive compatible and individually rational mechanisms, subject to budget balance. Because agents use quasi-linear utility functions, the expected social welfare is defined as the sum of all agents' expected utilities. By accumulating the expected utilities in (3.1) and (3.2) over all agents and substituting for the constraint of budget balance (3.10), the expected social welfare of the service allocation is calculated as follows:

$$
\begin{aligned}
&\sum_{i=1}^{N}\int_{0}^{\bar{\theta}} U_C(\theta_i) f(\theta_i)\, d\theta_i + \sum_{j=1}^{M}\int_{0}^{\bar{\sigma}} U_P(\sigma_j) h(\sigma_j)\, d\sigma_j \\
&\quad = \sum_{i=1}^{N}\int_{0}^{\bar{\theta}} \mathbb{E}_{\theta_{-i},\sigma}\Big[\sum_{j=1}^{M} v(\theta_i,\sigma_j) x_{ij}(\theta_i,\theta_{-i},\sigma) - t_C(\theta_i,\theta_{-i},\sigma)\Big] f(\theta_i)\, d\theta_i \\
&\quad\quad + \sum_{j=1}^{M}\int_{0}^{\bar{\sigma}} \mathbb{E}_{\theta,\sigma_{-j}}\Big[t_P(\theta,\sigma_j,\sigma_{-j}) \\
&\quad\quad\quad - \sum_{i=1}^{N} c(\theta_i,\sigma_j) x_{ij}(\theta,\sigma_j,\sigma_{-j})\Big] h(\sigma_j)\, d\sigma_j \\
&\quad = \mathbb{E}\Big[\sum_{i=1}^{N}\sum_{j=1}^{M} (v(\theta_i,\sigma_j) - c(\theta_i,\sigma_j)) x_{ij}(\theta,\sigma)\Big].
\end{aligned}
\tag{3.11}
$$

The mechanism's objective is to maximize the expression obtained in (3.11) over all agents. Therefore, the mechanism faces the optimization problem

$$
\max_{x_{ij}} \mathbb{E}\Big[\sum_{i=1}^{N}\sum_{j=1}^{M} (v(\theta_i,\sigma_j) - c(\theta_i,\sigma_j)) x_{ij}(\theta,\sigma)\Big], \tag{3.12}
$$

subject to

- incentive compatibility as defined in Section 3.3,
- individual rationality as defined in Section 3.4, and
- feasibility constraints as defined by

$$
0 \leq \sum_{j=1}^{M} x_{ij}(\theta,\sigma) \leq 1 \quad \forall a_i \qquad \text{and} \qquad 0 \leq \sum_{i=1}^{N} x_{ij}(\theta,\sigma) \leq 1 \quad \forall b_j. \tag{3.13}
$$

Notice that the budget balance constraint in Section 3.5 is always fulfilled in this setting because it has been incorporated into the derivation of the objective function (3.11).

Chapter 4

Allocation Mechanism

At the present time your plenty will
supply what they need, so that in turn
their plenty will supply what you need.
The goal is equality.

2 Corinthians 8:14, The Bible
(New International Version)

This chapter characterizes the economic properties of the mechanism based on the model introduced in Chapter 3 and defines the mechanism's allocation rule and expected monetary transfers. It then lists conditions sufficient for this mechanism to maximize the expected social welfare. Finally, a numerical example illustrates the allocation process for a market with uniformly distributed QoS.

4.1 Characterizing economic properties

As discussed in Section 3.6, ex post optimality cannot be achieved when incentive compatibility, individual rationality, and budget balance are required. Therefore, this research studies the associated second-best mechanism that maximizes the expected social welfare among all incentive compatible and individually rational mechanisms. Related literature in mechanism design theory, including Myerson and Satterthwaite

(1983), as well as Gresik and Satterthwaite (1989), have suggested using this ex ante performance measure instead of the corresponding ex post criterion. Similar to these approaches, the following sections define two functions that are crucial in constructing the second-best mechanism. Subsequently, incentive compatibility and individual rationality constraints are characterized for the given setting, followed by some results concerning budget balance.

4.1.1 Virtual valuation & virtual cost

In mechanism design theory, the concept of *virtual valuations* plays a central role when designing Bayesian optimal mechanisms. Let the *virtual valuation* of customer a_i be given by

$$\psi_C(\theta_i,\sigma_j) = v(\theta_i,\sigma_j) - \frac{1-F(\theta_i)}{f(\theta_i)}\frac{\partial v(\theta_i,\sigma_j)}{\partial \theta_i}. \tag{4.1}$$

The virtual valuation function of a_i is based on his actual valuation defined in Section 3.2.2, but subtracts from it a term that represents the *informational rent* of the agent. This informational rent represents the inevitable revenue loss caused by the private information of the agents. It can be interpreted as the surplus each customer must pay the mechanism for providing incentives to the agents to honestly declare their types (Börgers, 2015). Hence, the virtual valuation of a_i expresses the strategic behavior of a_i in an indirect mechanism implementation. Since $v(\theta_i,\sigma_j)$ is maximized for $\theta_i = \sigma_j$ and increasing in that maximal value, $\psi_C(\theta_i,\sigma_j)$ is strictly smaller than the actual valuation. Hence, customers have an incentive to understate their true valuation. They behave strategically in this fashion in order to influence their transaction prices.

Similar to the setting for customers, providers are endued with *virtual provision cost* functions. Virtual provision cost functions are based on the actual provision costs of providers but add to them the informational rent that providers are charged by the mechanism on top of their real provision cost. Insofar, each provider's informational rent can be viewed as a surplus that is withheld by the mechanism to provide adequate incentives for honest type declaration among all agents. The virtual provision cost of providers is defined by

$$\psi_P(\theta_i, \sigma_j) = c(\theta_i, \sigma_j) + \frac{H(\sigma_j)}{h(\sigma_j)} \frac{\partial c(\theta_i, \sigma_j)}{\partial \sigma_j}. \tag{4.2}$$

The virtual provision cost reflects the strategic behavior of providers in an indirect mechanism implementation. Since $c(\theta_i, \sigma_j)$ is minimized when $\sigma_j = \theta_i$ and increasing in that minimal value, the virtual cost is strictly greater than the actual provision cost. Thus, a provider has an incentive to overstate his true cost. In this way, providers try to raise their transaction prices for the offered service.

4.1.2 Incentive compatibility & individual rationality

As discussed in Section 2.2, the optimal allocation mechanism must satisfy two specific requirements from economic theory. First, the mechanism must ensure that no agent has an incentive to misreport his true type, assuming that all other agents report truthfully (incentive compatibility). Second, no agent can be forced to participate in the mechanism but must be given the liberty to withdraw from it at the interim stage; that is, when all agents know their own private information, while all the other agents have not yet reported their private information (individual rationality). The following Lemma characterizes the set of all incentive compatible and individually rational mechanisms for the allocation problem in this research.

Lemma 4.1. *Let $x_{ij}(\theta, \sigma)$ be the probability that provider b_j is allocated to customer a_i. Then transfer functions $t_C(\theta, \sigma)$ and $t_P(\theta, \sigma)$ exist such that $\{x_{ij}(\theta, \sigma), t_C(\theta, \sigma), t_P(\theta, \sigma)\}$ is incentive compatible and individually rational if* *and only if* $\mathbb{E}_{\theta_{-i},\sigma}[x_{ij}(\theta_i,\cdot)]$ *is non-decreasing,* $\mathbb{E}_{\theta,\sigma_{-j}}[x_{ij}(\cdot,\sigma_j)]$ *is non-increasing and*

$$\mathbb{E}\Big[\sum_{i=1}^{N}\sum_{j=1}^{M}(\psi_C(\theta_i,\sigma_j)-\psi_P(\theta_i,\sigma_j))x_{ij}(\theta,\sigma)\Big]\geq 0. \tag{4.3}$$

Proof. Suppose $\{x_{ij}(\theta,\sigma),t_C(\theta,\sigma),t_P(\theta,\sigma)\}$ is incentive compatible. Initially, the argument is derived for providers. For any quality pair $\hat{\sigma}\neq\sigma_j$ we must have

$$U_P(\sigma_j)\geq\mathbb{E}_{\theta,\sigma_{-j}}\Big[t_P(\theta,\hat{\sigma},\sigma_{-j})-\sum_{i=1}^{N}c(\theta_i,\sigma_j)x_{ij}(\theta,\hat{\sigma},\sigma_{-j})\Big]\quad\text{and} \tag{4.4}$$

$$U_P(\hat{\sigma})\geq\mathbb{E}_{\theta,\sigma_{-j}}\Big[t_P(\theta,\sigma_j,\sigma_{-j})-\sum_{i=1}^{N}c(\theta_i,\hat{\sigma})x_{ij}(\theta,\sigma_j,\sigma_{-j})\Big]. \tag{4.5}$$

These two inequalities imply that

$$\int_{\sigma_j}^{\hat{\sigma}}\mathbb{E}_{\theta,\sigma_{-j}}\Big[\sum_{i=1}^{N}\Big(x_{ij}(\theta,t,\sigma_{-j})-x_{ij}(\theta,\hat{\sigma},\sigma_{-j})\Big)\frac{\partial}{\partial t}c(\theta_i,t)\Big]dt\geq 0. \tag{4.6}$$

Therefore, if $\hat{\sigma}>\sigma_j$,

$$\mathbb{E}_{\theta,\sigma_{-j}}[x_{ij}(\theta,\sigma_j,\sigma_{-j})]\geq\mathbb{E}_{\theta,\sigma_{-j}}[x_{ij}(\theta,\hat{\sigma},\sigma_{-j})], \tag{4.7}$$

and thus $\mathbb{E}_{\theta,\sigma_{-j}}[x_{ij}(\cdot,\sigma_j)]$ is non-increasing. A similar argument for customers shows that $\mathbb{E}_{\theta_{-i},\sigma}[x_{ij}(\theta_i,\cdot)]$ must be non-decreasing. Corollary 1 in Milgrom and Segal (2002) provides expressions for the *indirect utility* of each agent in any incentive compatible mechanism:

$$U_C(\theta_i) = U_C(0) + \int_0^{\theta_i} \mathbb{E}_{\theta_{-i},\sigma}\Big[\sum_{j=1}^{M} \frac{\partial v(r,\sigma_j)}{\partial r} x_{ij}(r,\theta_{-i},\sigma)\Big] dr \quad \text{and} \tag{4.8}$$

$$U_P(\sigma_j) = U_P(\bar{\sigma}) + \int_{\sigma_j}^{\bar{\sigma}} \mathbb{E}_{\theta,\sigma_{-j}}\Big[\sum_{i=1}^{N} \frac{\partial c(\theta_i,r)}{\partial r} x_{ij}(\theta,r,\sigma_{-j})\Big] dr, \tag{4.9}$$

where $U_C(0)$ and $U_P(\bar{\sigma})$ are the expected utilities evaluated at the lower and upper quality bounds, respectively. By substituting the indirect utilities (4.8) and (4.9) into the sum of all agents' expected utilities given in (3.1) and (3.2), an alternative expression for the expected social welfare defined within the maximization problem in (3.12) is obtained:

$$\begin{aligned}
&\sum_{i=1}^{N} \int_0^{\bar{\theta}} U_C(\theta_i) f(\theta_i)\, d\theta_i + \sum_{j=1}^{M} \int_0^{\bar{\sigma}} U_P(\sigma_j) h(\sigma_j)\, d\sigma_j \\
&\quad = \sum_{i=1}^{N} \int_0^{\bar{\theta}} \Big(U_C(0) + \int_0^{\theta_i} \mathbb{E}_{\theta_{-i},\sigma}\Big[\sum_{j=1}^{M} \frac{\partial v(r,\sigma_j)}{\partial r} x_{ij}(r,\theta_{-i},\sigma)\Big] dr\Big) f(\theta_i)\, d\theta_i \\
&\qquad + \sum_{j=1}^{M} \int_0^{\bar{\sigma}} \Big(U_P(\bar{\sigma}) + \int_{\sigma_j}^{\bar{\sigma}} \mathbb{E}_{\theta,\sigma_{-j}}\Big[\sum_{i=1}^{N} \frac{\partial c(\theta_i,r)}{\partial r} x_{ij}(\theta,r,\sigma_{-j})\Big] dr\Big) h(\sigma_j)\, d\sigma_j.
\end{aligned} \tag{4.10}$$

Expression (4.10) is the expected social welfare expressed by the agents' *indirect* utilities. Therefore, it must equal the expected social welfare obtained by the agents' *direct* utilities in the maximand of (3.12). Equating these two expressions, followed by some basic algebraic manipulations, integration by parts, rearrangement and collection of similar terms yields

$$\sum_{i=1}^{N} U_C(0) + \sum_{j=1}^{M} U_P(\bar{\sigma}) = \mathbb{E}\Big[\sum_{i=1}^{N}\sum_{j=1}^{M} (\psi_C(\theta_i,\sigma_j) - \psi_P(\theta_i,\sigma_j)) x_{ij}(\theta,\sigma)\Big]. \tag{4.11}$$

Since individual rationality holds, $U_C(0) \geq 0$ and $U_P(\bar{\sigma}) \geq 0$ apply, which gives expression (4.3) in Lemma 4.1.

Suppose now that $x_{ij}(\theta, \sigma)$ satisfies (4.3) and that $\mathbb{E}_{\theta_{-i},\sigma}[x_{ij}(\theta_i, \cdot)]$ is non-decreasing and $\mathbb{E}_{\theta,\sigma_{-j}}[x_{ij}(\cdot, \sigma_j)]$ is non-increasing. Consider the following expected payments made by customer a_i

$$\begin{aligned} &\mathbb{E}_{\theta_{-i},\sigma}[t_C(\theta_i, \theta_{-i}, \sigma)] \\ &\quad = \mathbb{E}_{\theta_{-i},\sigma}\Big[\sum_{j=1}^{M} v(\theta_i, \sigma_j) x_{ij}(\theta_i, \theta_{-i}, \sigma) - \int_0^{\theta_i} \sum_{j=1}^{M} \frac{\partial v(r, \sigma_j)}{\partial r} x_{ij}(r, \theta_{-i}, \sigma) dr\Big] \end{aligned} \tag{4.12}$$

and the expected compensation received by provider b_j

$$\begin{aligned} &\mathbb{E}_{\theta,\sigma_{-j}}[t_P(\theta, \sigma_j, \sigma_{-j})] \\ &\quad = \mathbb{E}_{\theta,\sigma_{-j}}\Big[\sum_{i=1}^{N} c(\theta_i, \sigma_j) x_{ij}(\theta, \sigma_j, \sigma_{-j}) + \int_{\sigma_j}^{\bar{\sigma}} \sum_{i=1}^{N} \frac{\partial c(\theta_i, r)}{\partial r} x_{ij}(\theta, r, \sigma_{-j}) dr\Big]. \end{aligned} \tag{4.13}$$

These expected transfers are obtained by equating the direct and indirect expected utilities, as well as by setting the worst-off quality payoffs to zero; that is, $U_C(0) = U_P(\bar{\sigma}) = 0$.

It remains to be verified that these expected transfers satisfy incentive compatibility. To check incentive compatibility of (4.12), observe that

$$
\begin{aligned}
&U_C(\theta_i,\theta_{-i},\sigma) - U_C(\hat{\theta},\theta_{-i},\sigma) \\
&\quad = \mathbb{E}_{\theta_{-i},\sigma}\Big[\sum_{j=1}^{M} v(\theta_i,\sigma_j)(x_{ij}(\theta_i,\theta_{-i},\sigma) - x_{ij}(\hat{\theta},\theta_{-i},\sigma)) \\
&\qquad\quad - (t_C(\theta_i,\theta_{-i},\sigma) - t_C(\hat{\theta},\theta_{-i},\sigma))\Big] \\
&\quad = \mathbb{E}_{\theta_{-i},\sigma}\Big[\sum_{j=1}^{M} v(\theta_i,\sigma_j)\int_{\hat{\theta}}^{\theta_i}\frac{\partial}{\partial r}x_{ij}(r,\theta_{-i},\sigma)dr \\
&\qquad\quad - \Big(\sum_{j=1}^{M} v(\theta_i,\sigma_j)x_{ij}(\theta_i,\theta_{-i},\sigma) - v(\hat{\theta},\sigma_j)x_{ij}(\hat{\theta},\theta_{-i},\sigma) \\
&\qquad\quad + \int_{\theta_i}^{\hat{\theta}}\sum_{j=1}^{M}\frac{\partial v(r,\sigma_j)}{\partial r}x_{ij}(r,\theta_{-i},\sigma)dr\Big)\Big] \\
&\quad = \mathbb{E}_{\theta_{-i},\sigma}\Big[\sum_{j=1}^{M} v(\theta_i,\sigma_j)\int_{\hat{\theta}}^{\theta_i}\frac{\partial}{\partial r}x_{ij}(r,\theta_{-i},\sigma)dr \\
&\qquad\quad - \sum_{j=1}^{M}\int_{\hat{\theta}}^{\theta_i} v(r,\sigma_j)\frac{\partial}{\partial r}x_{ij}(r,\theta_{-i},\sigma)dr\Big] \\
&\quad = \mathbb{E}_{\theta_{-i},\sigma}\Big[\int_{\hat{\theta}}^{\theta_i}\sum_{j=1}^{M}(v(\theta_i,\sigma_j) - v(r,\sigma_j))\frac{\partial}{\partial r}x_{ij}(r,\theta_{-i},\sigma)dr\Big] \\
&\quad \geq 0.
\end{aligned}
\tag{4.14}
$$

The last inequality is a consequence of $\mathbb{E}_{\theta_{-i},\sigma}[x_{ij}(\theta_i,\cdot)]$ being non-decreasing. Recall that customer's valuation is increasing in the maximal value. Therefore, customer a_i would do better reporting θ_i instead of $\hat{\theta}$. The proof of incentive compatibility for providers is analogous.

Because (4.3) holds by assumption, the sum over all expected utilities evaluated at the lowest and highest qualities respectively must be non-negative. Further, equations (4.8) and (4.9) imply that $U_C(\theta_i)$ is increasing in θ_i and $U_P(\sigma_j)$ is decreasing in σ_j. Due to these monotonicity properties and because of (4.3), it suffices to verify individual rationality for customer's lowest desired quality $\theta_i = 0$ and provider's highest offered quality $\sigma_j = \bar{\sigma}$. This yields $U_C(\theta_i) \geq 0$ and $U_P(\sigma_j) \geq 0$. $\square$

4.1.3 Budget balance

As discussed in Section 2.2.6, the budget balance property is required to allow for distributed decision-making among the agents. Budget-balanced mechanisms do not depend on external sources of funds. Lemma 4.1 is crucial for constructing ex ante optimal mechanisms because it states that if the decision variables $x_{ij}(\theta,\sigma)$ satisfy the monotonicity properties, as well as inequality (4.3), then transfers $t_C(\theta,\sigma)$ and $t_P(\theta,\sigma)$ exist, such that $\{x_{ij}(\theta,\sigma), t_C(\theta,\sigma), t_P(\theta,\sigma)\}$ is an incentive compatible and individually rational mechanism. However, along with incentive compatibly and individual rationality, the budget balance property of the mechanism must still be verified. The next Lemma establishes that the mechanisms identified by Lemma 4.1 also satisfy budget balance.

Lemma 4.2. *Any incentive compatible, individually rational mechanism satisfies ex ante budget balance.*

Proof. To check ex ante budget balance, it must be shown that the net amount of payments made by customers and compensations received by providers never runs a deficit. Subtracting the sum of expected compensations from the sum of expected payments yields

$$\begin{aligned}
&\sum_{i=1}^{N}\int_0^{\bar{\theta}} \mathbb{E}_{\theta_{-i},\sigma}[t_C(\theta_i,\theta_{-i},\sigma)]f(\theta_i)d\theta_i - \sum_{j=1}^{M}\int_0^{\bar{\sigma}} \mathbb{E}_{\theta,\sigma_{-j}}[t_P(\theta,\sigma_j,\sigma_{-j})]h(\sigma_j)d\sigma_j \\
&= \mathbb{E}\Big[\sum_{i=1}^{N}\sum_{j=1}^{M} v(\theta_i,\sigma_j)x_{ij}(\theta_i,\theta_{-i},\sigma) - \sum_{i=1}^{N}\sum_{j=1}^{M}\frac{1-F(\theta_i)}{f(\theta_i)}\frac{\partial v(r,\sigma_j)}{\partial r}x_{ij}(r,\theta_{-i},\sigma)\Big] \\
&\quad - \mathbb{E}\Big[\sum_{j=1}^{M}\sum_{i=1}^{N} c(\theta_i,\sigma_j)x_{ij}(\theta,\sigma_j,\sigma_{-j}) + \sum_{j=1}^{M}\sum_{i=1}^{N}\frac{H(\sigma_j)}{h(\sigma_j)}\frac{\partial c(\theta_i,r)}{\partial r}x_{ij}(\theta,r,\sigma_{-j})\Big] \\
&= \mathbb{E}\Big[\sum_{i=1}^{N}\sum_{j=1}^{M}(\psi_C(\theta_i,\sigma_j)-\psi_P(\theta_i,\sigma_j))x_{ij}(\theta,\sigma)\Big] \geq 0. \qquad (4.15)
\end{aligned}$$

The inequality in line (4.15) is a consequence of (4.3) in any incentive compatible and individually rational mechanism. Hence, the net balance of what customers pay and what providers receive in the mechanism is greater than or equal to zero in expectation. Exact equality can always be achieved by subtracting a constant from customers' payments. □

4.2 Optimization

As discussed in Section 3.6, the ultimate objective of the allocation mechanism is to maximize the expected social welfare. Section 4.1.2 characterized the set of all allocation mechanisms that satisfy incentive compatibility and individual rationality (cf. Lemma 4.1). This section provides the definition of the direct revelation mechanism introduced in Section 3.2.4 in two steps. First, the allocation rule of the mechanism is specified. Then, a set of conditions are provided that are sufficient for the mechanism to maximize the expected social welfare among all incentive compatible and individually rational mechanisms.

4.2.1 Allocation rule & expected transfers

An allocation rule specifies the conditions for service exchange to take place between customers and providers and determines the expected monetary transfers. The definition of the allocation rule in this research requires some additional notation. Let the difference in virtual valuation and virtual provision cost be defined as

$$\psi(\theta_i, \sigma_j) = \psi_C(\theta_i, \sigma_j) - \psi_P(\theta_i, \sigma_j). \tag{4.16}$$

Using this definition, the *reserve functions* of customer a_i and provider b_j are respectively given by

$$R_C(\theta_i) = \begin{cases} \sup\{\sigma_j \in [0,\bar{\sigma}] : \psi(\theta_i,\sigma_j) \geq 0\} & \text{if} \quad \{\sigma_j : \psi(\theta_i,\sigma_j) \geq 0\} \neq \emptyset \\ 0 & \text{otherwise,} \end{cases} \tag{4.17}$$

$$R_P(\sigma_j) = \begin{cases} \inf\{\theta_i \in [0,\bar{\theta}] : \psi(\theta_i,\sigma_j) \geq 0\} & \text{if} \quad \{\theta_i : \psi(\theta_i,\sigma_j) \geq 0\} \neq \emptyset \\ \bar{\theta} & \text{otherwise.} \end{cases} \tag{4.18}$$

Reserve function $R_C(\theta_i)$ is used by customer a_i to identify the worst possible provider (i.e., the one with the highest offered quality) under the constraint $\psi(\theta_i,\sigma_j) \geq 0$, given his own desired quality θ_i. Reserve function $R_P(\sigma_j)$ is used by provider b_j to identify the worst possible customer (i.e., the one with the lowest desired quality) under the same constraint, given his own offered quality σ_j. These reserve functions correspond to the traders' reserve prices in private value auctions. A reserve price is the lowest price that a seller is willing to accept for selling the object (Krishna, 2002).

In addition, the lowest possible desired quality among all customers and, respectively, the highest possible offered quality among all providers are defined as

$$\underline{\theta}_i = \inf\{\theta_i \in [0,\bar{\theta}] : \{\sigma_j : \psi(\theta_i,\sigma_j) \geq 0\} \neq \emptyset\}\} \quad \text{and} \tag{4.19}$$

$$\bar{\sigma}_j = \sup\{\sigma_j \in [0,\bar{\sigma}] : \{\theta_i : \psi(\theta_i,\sigma_j) \geq 0\} \neq \emptyset\}\}. \tag{4.20}$$

The customer with desired quality $\underline{\theta}_i$ is the worst-off customer in the market, while the provider with offered quality $\bar{\sigma}_j$ is the worst-off provider. Let

$$\rho_\theta(\theta_i) = |\{\theta_k \in \theta : \theta_k \geq \theta_i\}| \tag{4.21}$$

denote the rank of desired quality θ_i within the vector of all customers' desired qualities $\theta = \{\theta_1, \ldots, \theta_N\}$. Define $\rho_\sigma(\sigma_i)$ similarly for providers. The quantity

$$w_k^C(\theta_i) = \frac{(N-1)!}{(N-k)!(k-1)!} F(\theta_i)^{N-k}(1-F(\theta_i))^{k-1} \tag{4.22}$$

gives the probability that customer's desired quality θ_i has rank k within vector θ. Define $w_k^P(\sigma_j)$ similarly for providers. Further, let $f_{(k)}(\cdot)$ and $h_{(k)}(\cdot)$ be the probability density functions of the k-th order statistic of customers and providers respectively. In statistics, the k-th order statistic of a sample is equal to the k-th smallest value of that sample (David and Nagaraja, 2003). The associated cumulative distribution functions are denoted by $F_{(k)}(\cdot)$ and $H_{(k)}(\cdot)$.

As pointed out in Section 3.2.4, the direct revelation mechanism in this research consists of the set of decision variables and the set of expected monetary transfers among the agents. The following definition summarizes these variables and specifies the direct revelation mechanism proposed in this research.

Definition 4.1. The *Truncated Positive Assortative Allocation (TPAA) mechanism* is defined by the allocation rule

$$x_{ij}(\theta, \sigma) = \begin{cases} 1 & \text{if} \quad \rho_\theta(\theta_i) = \rho_\sigma(\sigma_j) = k \quad \text{and} \quad \psi(\theta_i, \sigma_j) \geq 0, \\ 0 & \text{otherwise} \end{cases} \tag{4.23}$$

with expected payments of customers

$$\begin{aligned} T_C(\theta_i) = &\sum_{k=1}^{K} w_k^C(\theta_i) \int_0^{R_C(\theta_i)} v(\theta_i, r) h_{(k)}(r) dr \\ &- \int_{\underline{\theta}_i}^{\theta_i} \sum_{k=1}^{K} w_k^C(s) \int_0^{R_C(s)} \frac{\partial v(s,r)}{\partial s} h_{(k)}(r) dr ds \end{aligned} \tag{4.24}$$

and expected compensations for providers

$$T_P(\sigma_j) = \sum_{k=1}^{K} w_k^P(\sigma_j) \int_{R_P(\sigma_j)}^{\bar{\theta}} c(r,\sigma_j) f_{(k)}(r) dr \\ + \int_{\sigma_j}^{\bar{\sigma}_j} \sum_{k=1}^{K} w_k^P(s) \int_{R_P(s)}^{\bar{\theta}} \frac{\partial c(r,s)}{\partial s} f_{(k)}(r) dr ds. \tag{4.25}$$

The allocation rule of the TPAA mechanism captures two distinct features. First, a service is allocated from a provider to a customer if and only if the rank of a_i's desired quality is equal to the rank of b_j's offered quality. This means that the service from the provider with the highest offered quality is allocated to the customer with the highest desired quality, the second-highest, and so on. Such mechanisms are *positively assortative* (Shimer and Smith, 2000). Second, service allocation from a provider to a customer takes place if and only if the difference in a_i's virtual valuation and b_j's virtual provision cost is positive.

Notice that the associated transfers are the same as the expected transfers (4.12) and (4.13) derived in the proof of Lemma 4.1 but evaluated at the allocation rule (4.23). These transfers are similar to the payments in other mechanism design settings. A customer's expected payment is his expected surplus less his cost that incurs due to information revelation. Similarly, a provider's expected compensation equals his expected surplus plus a term he receives for revealing his information.

4.2.2 Maximizing expected social welfare

This section shows that the TPAA mechanism in Definition 4.1 maximizes the expected social welfare among all incentive compatible and individually rational mechanisms characterized in Lemma 4.1. The first argument demonstrates that the monotonicity properties required in Lemma 4.1 are satisfied. Then, the main result is stated in terms of Theorem 4.1.

Taking the partial derivative of the expected allocation probability and evaluating this variable at the allocation rule (4.23) for customers gives

$$\begin{aligned}
\frac{\partial}{\partial\theta_i}\mathbb{E}_{\theta_{-i},\sigma}[x_{ij}(\theta_i,\cdot)] &= \frac{\partial}{\partial\theta_i}\sum_{k=1}^{K} w_k^C(\theta_i)\int_0^{R_C(\theta_i)} h_{(k)}(r)dr \\
&= \frac{\partial}{\partial\theta_i}\sum_{k=1}^{K} w_k^C(\theta_i)H_{(k)}(R_C(\theta_i)) \\
&= \sum_{k=1}^{K}\frac{\partial w_k^C(\theta_i)}{\partial\theta_i}H_{(k)}(R_C(\theta_i)) + \sum_{k=1}^{K} w_k^C(\theta_i)h_{(k)}(R_C(\theta_i))R_C'(\theta_i).
\end{aligned} \tag{4.26}$$

Observe that the second summand in (4.26) is always positive because $R_C(\theta_i)$ is increasing. For the first summand the same argument as in Johnson (2011) is applied: Fix index k^*. Then increase all terms above $k > k^*$ where $\partial w_k^C(\theta_i)/\partial\theta_i$ is negative, and decrease all terms below $k < k^*$ with $\partial w_k^C(\theta_i)/\partial\theta_i$ being positive. By the binomial theorems (Abramowitz and Stegun, 1972), the following identities hold:

$$\sum_{k=1}^{K} w_k^C(\theta_i) = 1 \qquad \text{and} \qquad \sum_{k=1}^{K}\frac{\partial w_k^C(\theta_i)}{\partial\theta_i} = 0. \tag{4.27}$$

Thus, equation (4.26) can be estimated as follows:

$$\begin{aligned}\frac{\partial}{\partial\theta_i}\mathbb{E}_{\theta_{-i},\sigma}[x_{ij}(\theta_i,\cdot)] &> \sum_{k=1}^{K}\frac{\partial w_k^C(\theta_i)}{\partial\theta_i}H_{(k)}(R_C(\theta_i)) \\ &> H_{(k^*)}(R_C(\theta_i))\sum_{k=1}^{K}\frac{\partial w_k^C(\theta_i)}{\partial\theta_i} \\ &= 0, \qquad (4.28)\end{aligned}$$

where the last line follows from the first identity in (4.27). Since the derivative is positive, $\mathbb{E}_{\theta_{-i},\sigma}[x_{ij}(\theta_i,\cdot)]$ is increasing in θ_i.

For providers, allocation rule (4.23) yields

$$\begin{aligned}\frac{\partial}{\partial\sigma_j}\mathbb{E}_{\theta,\sigma_{-j}}[x_{ij}(\cdot,\sigma_j)] &= \frac{\partial}{\partial\sigma_j}\sum_{k=1}^{K}w_k^P(\sigma_j)\int_{R_P(\sigma_j)}^{\bar\theta}f_{(k)}(r)dr \\ &= \frac{\partial}{\partial\sigma_j}\sum_{k=1}^{K}w_k^P(\sigma_j)(1-F_{(k)}(R_P(\sigma_j))) \\ &= -\Big(\sum_{k=1}^{K}\frac{\partial w_k^P(\sigma_j)}{\partial\sigma_j}F_{(k)}(R_P(\sigma_j)) \\ &\qquad + \sum_{k=1}^{K}w_k^P(\sigma_j)f_{(k)}(R_P(\sigma_j))R_P'(\sigma_j)\Big). \qquad (4.29)\end{aligned}$$

Notice that $R_P(\sigma_j)$ is increasing. Labeling the indices k as k^* and using a similar argument as above shows that the expression in the outer brackets in (4.29) is always positive. Therefore, the whole term is negative, and thus $\mathbb{E}_{\theta,\sigma_{-j}}[x_{ij}(\cdot,\sigma_j)]$ is decreasing in σ_j. This proves that the monotonicity properties of Lemma 4.1 are satisfied. In addition, the allocation rule also fulfills constraint (4.3) because the rule is positive assortative as long as the difference between virtual valuation and virtual cost is positive.

The following theorem provides sufficient conditions for the TPAA mechanism to maximize the expected social welfare among all incentive compatible and individually rational mechanisms.

Theorem 4.1. *Let customer a_i's valuation $v(\theta_i, \sigma_j)$ be maximized at $\theta_i = \sigma_j$ and increasing in its maximum, supermodular, and concave in θ_i. Further, let*

$$\begin{aligned} &\frac{\partial^3 v(\theta_i, \sigma_j)}{\partial \theta_i^2 \partial \sigma_j} \geq 0, \\ &\frac{f(\theta_i)}{1 - F(\theta_i)} \geq \frac{\partial}{\partial \theta_i} \log\left(\frac{\partial v(\theta_i, \sigma_j)}{\partial \sigma_j}\right). \end{aligned} \tag{4.30}$$

On the supply side, let provider b_j's provision cost $c(\theta_i, \sigma_j)$ be minimized at $\sigma_j = \theta_i$ and increasing in its minimum, submodular, and convex in σ_j. Further, let

$$\begin{aligned} &\frac{\partial^3 c(\theta_i, \sigma_j)}{\partial \theta_i \partial \sigma_j^2} \leq 0, \\ &\frac{h(\sigma_j)}{H(\sigma_j)} \geq \frac{\partial}{\partial \sigma_j} \log\left(\frac{\partial c(\theta_i, \sigma_j)}{\partial \theta_i}\right). \end{aligned} \tag{4.31}$$

Let the distribution functions $F(\cdot)$ and $H(\cdot)$ be log-concave. Then the TPAA mechanism maximizes the expected social welfare among all incentive compatible and individually rational mechanisms. Further, budget balance is satisfied.

Proof. The inner term of the maximand in (3.12) is the difference between a_i's valuation and b_j's cost. Because $v(\theta_i, \sigma_j)$ is supermodular and $c(\theta_i, \sigma_j)$ is submodular (i.e., $-c(\theta_i, \sigma_j)$ is supermodular), their difference as a linear combination is again supermodular. If the production function is supermodular, the optimal match function is positively assortative (Shimer and Smith, 2000). By Lemma 4.1, the TPAA mechanism is incentive compatible and individually rational if and only if the monotonicity properties are satisfied. These are satisfied if the reserve functions $R_C(\theta_i)$

and $R_P(\sigma_j)$ are increasing. The hazard rate (4.30) for customers guarantees that a higher desired quality θ_i unambiguously results in a better lottery over the offered qualities of providers. Similarly, the hazard rate (4.31) on the provider side ensures that a higher offered quality σ_j unambiguously entails a lower likelihood to find a customer. Therefore, the reserve functions are increasing. Lemma 4.2 establishes budget balance. □

The notions of supermodularity and submodularity used in Theoreom 4.1 receive the following interpretation. In economic theory, the supermodularity property of a function represents the economic concept of complementary inputs (Milgrom and Roberts, 1990). Loosely speaking, a function is supermodular if it has "increasing differences" (Chambers and Echenique, 2009). In the light of Theoreom 4.1, the marginal utility of a customer with desired quality θ_i increases in the offered quality σ_j of a provider. This interpretation is intuitive because customers obtain higher utilities from matching with high-quality providers. The reverse explanation applies to the submodularity of providers' provision costs. For twice continuously differentiable functions, the cross-partial derivatives of supermodular (submodular) functions are non-negative (non-positive).

4.2.3 Numerical example

This section presents a numerical example to illustrate the design of second-best mechanisms that satisfy incentive compatibility, individual rationality, and budget balance. First, valuation and cost are defined based on the assumptions made in Theoreom 4.1. Then, the relationship between each agent's actual and virtual reservation value is discussed. Finally, the example depicts the allocation boundaries of the second-best mechanism in comparison to the associated first-best boundaries.

Suppose each customer's desired quality θ_i and each provider's offered quality σ_j are identically and uniformly distributed over the unit interval. All customers have valuations of

$$v(\theta_i, \sigma_j) = 1 + \sqrt{\theta_i} - (\theta_i - \sigma_j)^2 \tag{4.32}$$

and all providers have provision costs of

$$c(\theta_i, \sigma_j) = \sigma_j^2 + (\theta_i - \sigma_j)^2. \tag{4.33}$$

These functions satisfy the requirements in Theorem 4.1. In particular, the valuation $v(\theta_i, \sigma_j)$ is *supermodular* because

$$\frac{\partial^2}{\partial \theta_i \partial \sigma_j} v(\theta_i, \sigma_j) = 2 \geq 0, \tag{4.34}$$

and the provision cost $c(\theta_i, \sigma_j)$ is *submodular* because

$$\frac{\partial^2}{\partial \theta_i \partial \sigma_j} c(\theta_i, \sigma_j) = -2 \leq 0. \tag{4.35}$$

Since all private information is independently drawn from the uniform distribution, the densities are $f(\theta_i) = h(\sigma_j) = 1$, with cumulative distributions $F(\theta_i) = \theta_i$ and $H(\sigma_j) = \sigma_j$ over the interval $[0, 1]$. Given these functions, the virtual valuation of customers evaluates to

$$\psi^C(\theta_i, \sigma_j) = v(\theta_i, \sigma_j) - (1 - \theta_i)(2\sigma_j - 2\theta_i + \frac{1}{2\sqrt{\theta_i}}), \tag{4.36}$$

while the virtual cost of providers is

$$\psi^P(\theta_i, \sigma_j) = c(\theta_i, \sigma_j) + \sigma_j(4\sigma_j - 2\theta_i). \tag{4.37}$$

Figures 4.1 and 4.2 illustrate the relationship between actual and virtual reservation values of the agents. The fact that a customer's virtual valuation is less than or equal to his actual valuation indicates that he will in general have an incentive to understate his desired quality in order to favorably influence the transaction price. The customer with the strongest incentive to understate his desired quality is the one with the highest desired quality. However, since this customer will generate more social welfare (i.e., gains from trade) than a customer with a low desired quality, an efficient mechanism will match the highest quality customer with a provider of a lower actual cost whenever doing so is efficient. Hence, the virtual valuation of the highest quality customer is equal to his actual valuation. To ensure that this customer reports his desired quality truthfully, the mechanism must make lower reports unprofitable. It does so by lowering the probability that a customer with a lower desired quality will be matched at a lower probability. This reduced probability is inefficient, but the loss in social welfare is smaller from reducing the probability of trade for customers with low desired qualities than with high desired qualities. The reverse explanation applies to the pattern observed for providers. Hence, this example illustrates what the revelation principle does: It embeds the strategic behavior of the agents in the mechanism so that in equilibrium the agents no longer engage in strategic misrepresentation.

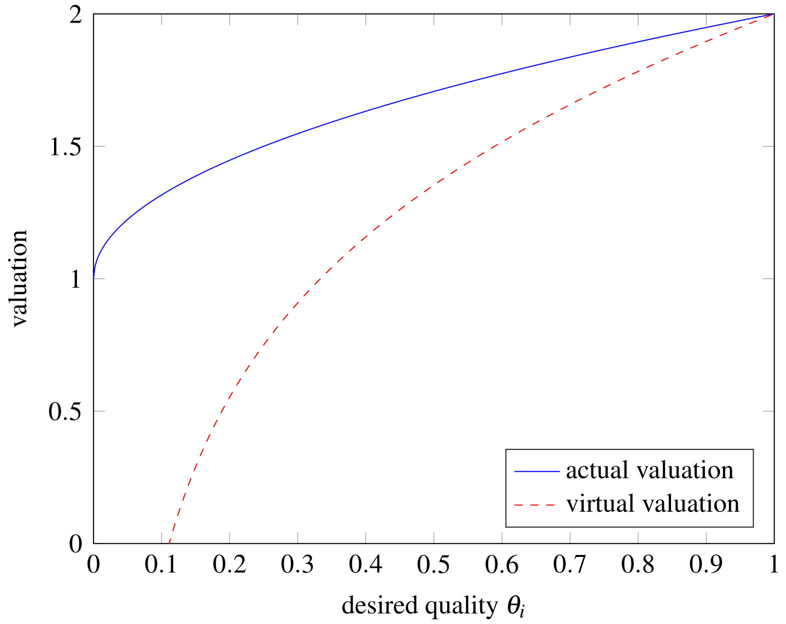

Figure 4.1: Actual and virtual valuation of customers.

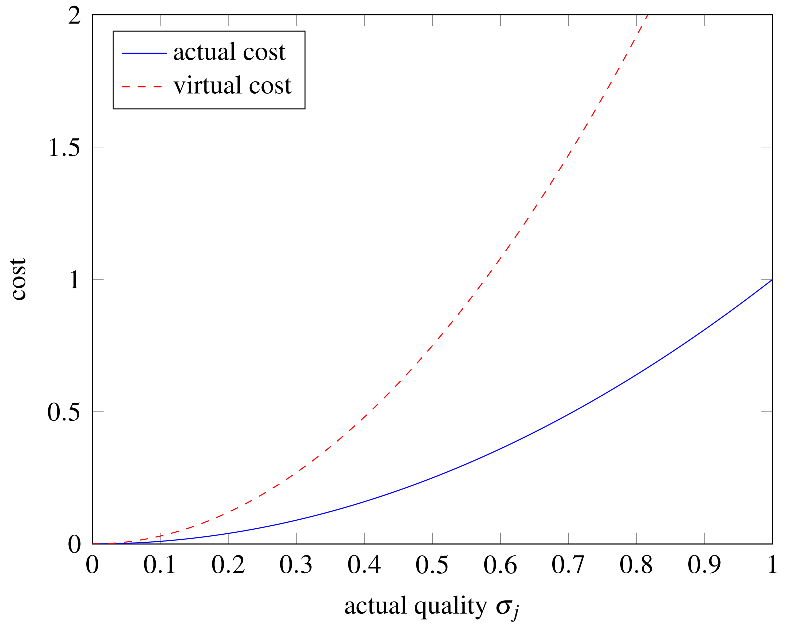

Figure 4.2: Actual and virtual provision cost of providers.

The mechanism arranges an allocation if and only if the difference in virtual valuation and virtual cost is non-negative. Figure 4.3 shows the boundaries of the first-best and second-best mechanism for the case of uniformly distributed qualities. The shaded area in solid green contains the eligible quality pairs within the domain $[0,1] \times [0,1]$ that warrant service allocation in the first-best mechanism. The dotted region within the solid area contains all quality pairs that are allocated in the second-best mechanism. The significantly smaller area in the second-best case illustrates the distortion due to the presence of privately known qualities on either side of the market.

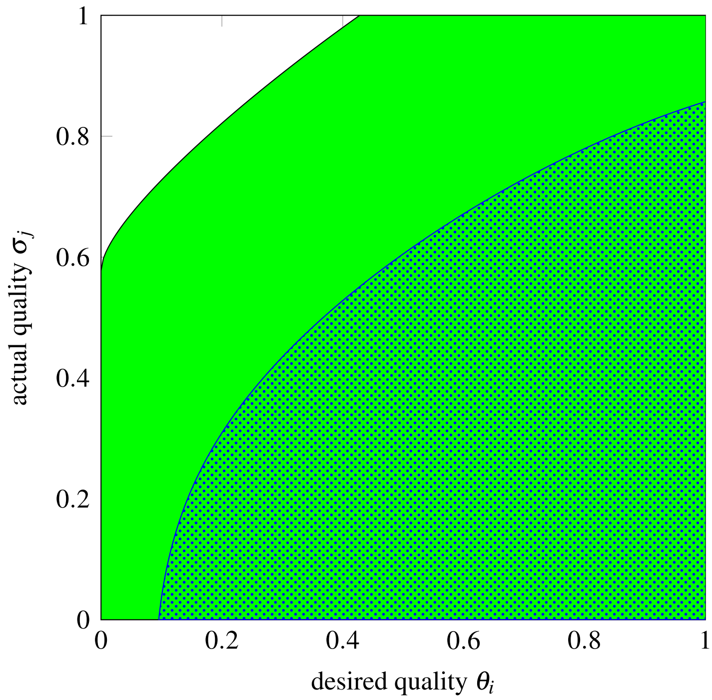

Figure 4.3: The shaded area (solid green) contains all eligible, *uniformly distributed* quality pairs that warrant service allocation in the first-best mechanism. The dotted area contains the corresponding quality pairs in the second-best mechanism.

Since the TPAA mechanism is positively assortative, it converges to deterministic quantile matching (Johnson, 2011). Therefore, in the limit as $K \to \infty$, the empirical quantile function is asymptotically equivalent to the k-th order statistic (van der Vaart, 2000). This means that $\sigma \to H^{-1}(F(\theta))$ in the allocation rule, and therefore $\sigma \to \theta$ for uniformly distributed θ and σ on $[0,1]$. Hence, substituting $\sigma = \theta$ into the allocation surplus yields $v(\theta,\theta) - c(\theta,\theta) = \sqrt{\theta} - \theta^2 + 1$, representing the expected match surplus generated by an allocation as $K \to \infty$ for uniformly distributed qualities. Figure 4.4 shows this surplus, while the two dashed lines indicate the lower and upper bound of the exclusion condition $\psi(\theta, \theta) \geq 0$. Solving this inequality gives the lower bound $a = 0.1157$ and the upper bound $b = 0.7604$ so that in the limit, all qualities less than a and greater than b are not allocated by the mechanism.

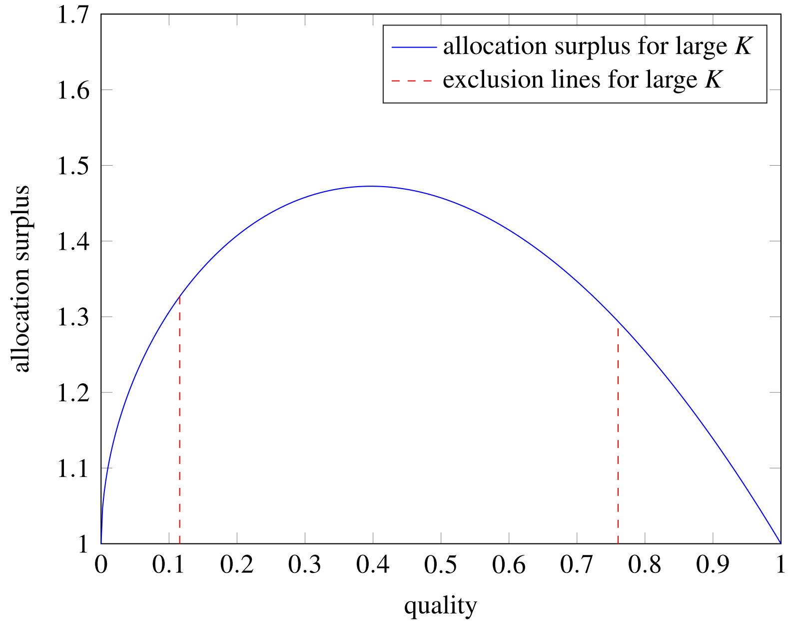

Figure 4.4: The solid curve shows the average allocation surplus for $K \to \infty$, while only the quality pairs between the dashed lines are matched under the uniform distribution.

For this rather simple example, the asymptotic efficiency of the second-best mechanism can be calculated analytically. As discussed in Section 2.1.2.4, this efficiency is given by the ratio between the outcome of the second-best mechanism and the outcome of the first-best mechanism if it were to exist. The outcome of the second-best mechanism must consider the lower and upper bound of the exclusion condition arising from incentive compatibility and individual rationality constraints. The outcome of the first-best mechanism is simply the average match surplus that would be generated if all QoS information were publicly known. Therefore, the asymptotic efficiency of the mechanism as $K \to \infty$ is given by

$$\frac{\displaystyle\int_a^b \Big(v(t,t) - c(t,t)\Big)\,dt}{\displaystyle\int_0^1 \Big(v(t,t) - c(t,t)\Big)\,dt} \approx 0.6859, \tag{4.38}$$

since $f(t) = 1$ for the uniform distribution. Thus, the TPAA mechanism's efficiency for uniformly distributed qualities over the unit interval is 68.59%. The experimental evaluation in Chapter 5 confirms this result (cf. convergence asymptote in Figure 5.3).

Chapter 5

Experimental Evaluation

Test all things; hold fast what is good.

1 Thessalonians 5:21, The Bible
(New King James Version)

This chapter reports on the experimental evaluation of the allocation mechanism for electronic services. The experimental evaluation was executed by simulation with artificial data. This chapter presents the experimental setup used for the simulation study, the results obtained through the simulation, and discusses the economic insights that can be inferred from these results.

5.1 Simulation setup

The simulation study in this research compares the outcome of the second-best mechanism to the one achieved by the associated first-best mechanism if it were to exist (cf. Section 2.1.2.4). This study considers a set of experiments that vary in the probability distribution of the private information as well as in the number of customers and providers in the market. All experiments are executed by implementing the allocation algorithm using MATLAB (2015). This section describes the different experiments and the allocation algorithm.

5.1.1 Experiments

The simulation study considers three distinct experiments executed by simulation with artificial data. These experiments differ in the choice of the underlying probability distributions for the QoS desired by customers and the QoS offered by providers. The experiments include:

- **Experiment 1 (symmetric, uniform):** The QoS realizations desired by customers and offered by providers in the market are independently drawn from the *uniform* distribution over the unit interval $[0,1]$; that is, $\theta_i \sim \sigma_j \sim U(0,1)$ for all a_i, b_j. Thus, demand and supply are symmetric.

- **Experiment 2 (symmetric, normal):** The QoS values of customers and providers are independently drawn from the *normal* distribution truncated to the unit interval $[0,1]$ with a mean of $\mu = 0.5$ and a standard deviation of $\sigma = 0.1$; that is, $\theta_i \sim \sigma_j \sim N(0.5, 0.01)$ for all a_i, b_j. Again, demand and supply are symmetric.

- **Experiment 3 (asymmetric, beta):** Unlike Experiments 1 and 2, this experiment assumes *asymmetric* demand and supply. The QoS desired by customers is independently drawn from the beta distribution with $\alpha = 3$ and $\beta = 2$; that is, $\theta_i \sim Beta(3,2)$ for all a_i. In contrast, the QoS realizations actually offered by providers are independently drawn from the beta distribution with $\alpha = 1$ and $\beta = 3$; that is, $\sigma_j \sim Beta(1,3)$ for all b_j. On such markets, customers tend to desire high QoS. The probability density function for $Beta(3,2)$ has a local maximum at $2/3$. That is, the relative likelihood of a customer to desire a QoS of less than $2/3$ increases monotonically. For QoS values above $2/3$, however, the relative likelihood decreases because customers anticipate higher prices for such high-quality services. On the supply side, providers are more likely to offer low

QoS since increased quality entails higher provision costs. Hence, the relative likelihood of providers to offer high QoS decreases monotonically. The probability density functions of customers and providers are shown in Figure 5.1.

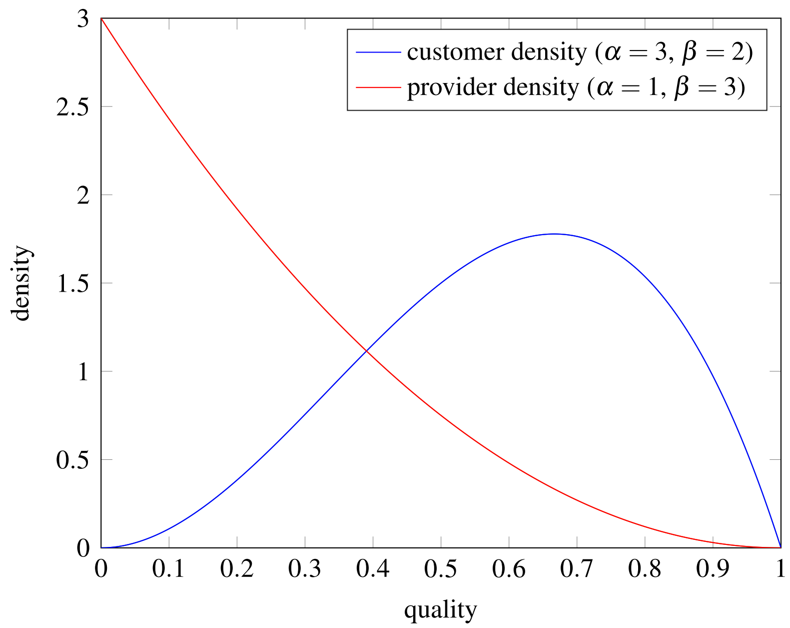

Figure 5.1: Density functions of beta-distributed customers and providers.

Because the simulation study evaluates the TPAA mechanism, it must be ensured that the conditions on the underlying probability distributions required in Theorem 4.1 are satisfied. First, the uniform, normal, and beta distribution must be log-concave. This requirement is fulfilled for all three distributions, given the parameters used in this study (Bagnoli and Bergstrom, 1989). Second, the hazard bounds stated in inequalities (4.30) and (4.31) must be satisfied such that the mechanism's reserve functions defined in equation (4.18) are increasing. The second-best boundaries depicted in figures 4.3 (uniform), 5.11 (normal), and 5.12 (beta) show that the reserve functions of both customers and providers are indeed increasing. Therefore, Theorem 4.1 is applicable for the given parameter settings.

For assessing different market structures, each of the three experiments is executed in two particular settings, namely *fixed* and *increasing*. Setting 1 (fixed) assumes a fixed number of providers in the market, while the number of customers increases. In this setting, there are 10, 50, and 100 providers that offer their services to a growing number of customers. Setting 2 (increasing) reflects a market in which the number of customers and the number of providers both increase at an equal rate. Thus, the two settings include:

- **Setting 1 (fixed):** The number of providers is fixed to 10, 50, and 100.
- **Setting 2 (increasing):** The number of providers is increasing.

Table 5.1 provides an overview of all three experiments executed in this simulation study.

Table 5.1: Experiments overview.

Experiment	Num. of cust. N	Num. of prov. M	Distribution
1	increasing	fixed	uniform symmetric
		increasing	
2	increasing	fixed	normal symmetric
		increasing	
3	increasing	fixed	beta asymmetric
		increasing	

The deterministic and the random experiment parameters for the three experiments are summarized in Tables 5.2 and 5.3. For every permutation of deterministic and random parameters, an experiment instance is generated. Each experiment is repeated 10^5 times and average values are calculated. For each repetition of the experiment, the QoS desired by customers and the QoS offered by providers are freshly and independently drawn from the respective probability distribution.

Table 5.2: Deterministic experiment parameters.

	Parameter	Range	Distribution
Fixed	N	$\{2,4,\ldots,98,100\}$	uniform, normal, and beta
	M	$\{10,50,100\}$	
Increasing	$K=\min\{N,M\}$	$\{10,50,100,\ldots,1000\}$	uniform
		$\{10,20,\ldots,200\}$	normal and beta

Table 5.3: Random experiment parameters.

	Probability distribution			
Parameter	**Exp. 1**	**Exp. 2**	**Exp. 3**	**Range**
Customer desired QoS θ_i	$U(0,1)$	$N(0.5,0.01)$	$Beta(3,2)$	$[0,1]$
Provider offered QoS σ_j	$U(0,1)$	$N(0.5,0.01)$	$Beta(1,3)$	$[0,1]$

The experimental setup is based on the setting used in the numerical example presented in Section 4.2.3. Thus, the valuation function of each customer is given by

$$v(\theta_i,\sigma_j) = 1+\sqrt{\theta_i}-(\theta_i-\sigma_j)^2, \tag{5.1}$$

and all providers have provision costs of

$$c(\theta_i,\sigma_j) = \sigma_j^2+(\theta_i-\sigma_j)^2. \tag{5.2}$$

As discussed in Section 4.2.3, these functions satisfy the requirements in Theorem 4.1.

5.1.2 Algorithm

For the purpose of the simulation study, the following algorithm is implemented using MATLAB (2015). The algorithm is repeated 10^5 times and average values are calculated. The steps for each experiment are defined as follows:

1. Desired and offered quality realizations are drawn independently from the respective random distributions.

2. Desired and offered quality realizations are sorted positive assortatively.

3. The allocation surplus is calculated as $v(\theta_i,\sigma_j)-c(\theta_i,\sigma_j)$.

4. The allocation constraint is calculated as $\psi_C(\theta_i, \sigma_j) - \psi_C(\theta_i, \sigma_j)$.

5. The sum of all allocation surpluses for the first-best mechanism (without constraint) as well as for the second-best mechanism (with constraint) are accumulated.

6. The efficiency of the mechanism is calculated as the ratio between the mean across all second-best outcomes and the mean across all first-best outcomes.

Algorithm 1 depicts the implementation of the TPAA mechanism in pseudocode. The input parameters of the algorithm are two vectors: the N-vector of all customers' desired quality values and the M-vector of all providers' offered quality values. Once both vectors are sorted in descending order, the algorithm calculates the value of each allocation (i.e., the match surplus). If the surplus is non-negative, the ex post welfare is incremented. Next, the constraint for incentive compatibility and individual rationality derived in Lemma 4.1 is calculated. If it is non-negative, the value for ex ante welfare is incremented. Finally, the efficiency of the mechanism is calculated as the ratio between the second-best outcome and the first-best outcome.

```
input  : Customers' desired quality vector θ = (θ_1, ..., θ_N) and providers'
         offered quality vector σ = (σ_1, ..., σ_M).
output : The efficiency of the mechanism.

welfare_expost ← 0;
welfare_exante ← 0;
surplus ← 0;
θ ← sortDescending(θ);
σ ← sortDescending(σ);
K ← min{N,M};
;
for i ← 1 to K do
    surplus ← v(θ(i),σ(i)) − c(θ(i),σ(i));
    if surplus > 0 then
        welfare_expost ← welfare_expost + surplus;
        constraint ← ψ^C(θ(i),σ(i)) − ψ^P(θ(i),σ(i));
        if constraint > 0 then
            welfare_exante ← welfare_exante + surplus;
        end
        ;
    end
    ;
end
;
efficiency ← welfare_exante/welfare_expost
```

Algorithm 1: Direct mechanism implementation.

As discussed in Section 2.1.2.4, the algorithm calculates the outcomes for both the first-best and the second-best mechanism. In the first-best case, the mechanism assumes the existence of *complete information* and thus arranges an allocation if and only if a customer's valuation exceeds a provider's provision cost. Since complete information is available to the agent, the mechanism does not need to consider the incentive compatibility and individual rationality constraints. The expected social welfare achieved by the first-best mechanism is then set as benchmark and compared to the outcome produced by the second-best mechanism under private information. Subsequently, the efficiency of the TPAA mechanism is computed as the ratio between the outcome of the first-best and the outcome of the second-best mechanism for each permutation of deterministic parameters. Finally, this simulation study estimates the behavior of the asymptotic efficiency for large markets.

5.2 Results

This section presents the results achieved by the simulation study. It reports the results obtained in Experiment 1 with uniform, symmetrically distributed QoS, in Experiment 2 with normally, symmetrically distributed QoS, and Experiment 3 with asymmetric, beta-distributed QoS.

5.2.1 Uniformly distributed, symmetric quality

In the sections that follow, the QoS desired by customers and the QoS offered by providers is drawn from the uniform probability distribution on the unit interval. All agents are assumed to be symmetric. First, the results for a fixed number of providers are presented, followed by the results for an increasing number of providers.

5.2.1.1 Fixed number of providers

Table 5.4 shows the results obtained by the TPAA mechanism for a market with an increasing number of customers and the number of providers fixed to 10, 50, and 100. All agents are symmetric uniform; that is, all QoS realizations are drawn from the uniform distribution over the unit interval. The columns labeled "$M = 10$", "$M = 50$" and "$M = 100$" show the efficiency of the TPAA mechanism with 10, 50, and 100 providers, while the number of customers increased from 2 to 100. The efficiency is the proportion of the *ex ante* expected social welfare relative to the social welfare that an *ex post* efficient mechanism would achieve if it were to exist (cf. Section 2.1.2.4).

Table 5.4: Efficiency of the TPAA mechanism for *uniformly* distributed QoS with the number of providers fixed to 10, 50, and 100.

$N = \|\mathscr{A}_C\|$	$M = 10$	$M = 50$	$M = 100$
2	0.1312	0.0001	0.0000
10	0.5921	0.0011	0.0000
20	0.8312	0.0139	0.0000
30	0.8450	0.0917	0.0001
40	0.8479	0.3371	0.0010
50	0.8477	0.6288	0.0071
60	0.8482	0.8085	0.0351
70	0.8473	0.8314	0.1237
80	0.8470	0.8419	0.2994
90	0.8464	0.8476	0.5045
100	0.8464	0.8508	0.6536

Figure 5.2 depicts the content of Table 5.4 graphically. It shows the average behavior of the mechanism's efficiency as a function of the number of customers for uniformly distributed desired and offered QoS realizations. The number of providers on the market is fixed to 10, 50, and 100. For 10 providers, the efficiency is at its minimum of 0.1312, followed by an efficiency increase as the number of customers grows. At about 60 customers, the efficiency reaches its maximum at 0.8482. Then, the efficiency decreases at a very slow rate and arrives at 0.8464 for 100 customers. Further experiments with larger numbers of customers (and fixed provider numbers) show that the efficiency values scatter around the asymptote at 0.8397. This asymptotic efficiency of 0.8397 can be derived analytically as follows. When the number of customers increases to infinity, the corresponding limit qualities that are matched with providers' qualities will converge to 1. This convergence property is due to the fact that the TPAA mechanism arranges allocations positive assortatively with

$K = \min\{N, M\}$ agents on either side of the market (cf. Section 4.2.3). Therefore, in the limit with $N \to \infty$ and fixed M, the efficiency is given by

$$\frac{\int_a^b \Big(v(1,t) - c(1,t)\Big)\, dt}{\int_0^1 \Big(v(1,t) - c(1,t)\Big)\, dt} \approx 0.8397, \tag{5.3}$$

where the density is $f(t) = 1$ for the uniform distribution. The bounds of integration $a = 0$ and $b = 0.8571$ are the solutions to $\psi_C(1,t) - \psi_P(1,t) = 0$. These bounds reflect the exclusion boundaries of the second-best mechanism as depicted in the right hand side of Figure 4.3. Hence, in the limit, all quality values of providers that exceed $b = 0.8571$ are excluded by the mechanism.

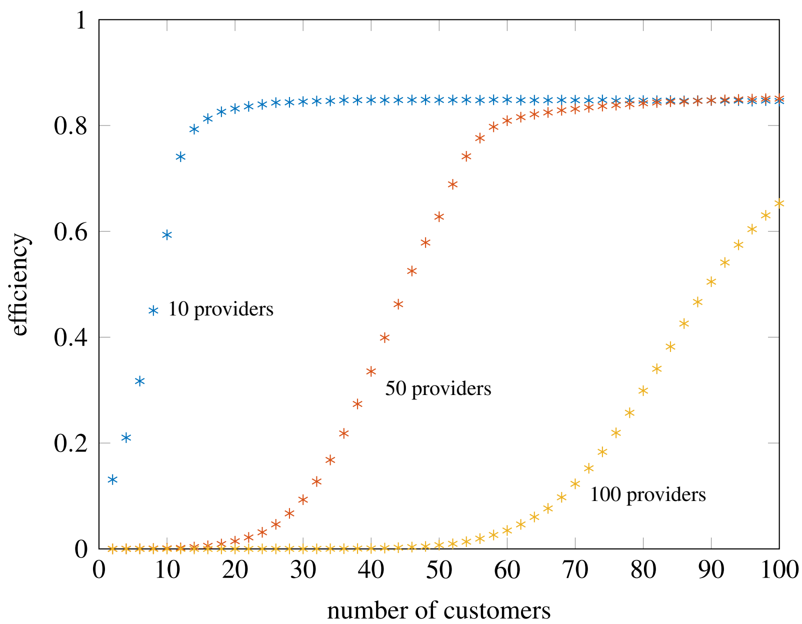

Figure 5.2: Efficiency of the TPAA mechanism with symmetric agents and QoS realizations drawn from the *uniform distribution* over the unit interval with 10, 50, and 100 providers.

When 50 providers offer their services to less than 20 customers, the efficiency vanishes completely. For 20 to 70 customers the efficiency increases. For 80 and more customers, the efficiency exceeds the asymptote, reaching a maximum at 0.8508 for 100 customers. Similar to the experiment with 10 providers, further experiments indicate that the efficiency for customer numbers beyond 100 scatters around the asymptote at 0.8397. With 100 providers on the market, the efficiency is zero for less than 50 customers. It then increases until the same limit of 0.8397 is reached (not shown in the graph).

5.2.1.2 Increasing number of providers

Table 5.5 shows the efficiency of the TPAA mechanism when the number of customers and the number of providers is increased equally. The column labeled "Second-best welfare" contains the expected social welfare that the TPAA mechanism generates. The column labeled "First-best welfare" contains the expected social welfare that the corresponding ex post efficient mechanism would achieve if such a mechanism were to exist. The "Efficiency" column represents the proportion of the expected social welfare that the ex ante efficient mechanism generates relative to the expected social welfare that an ex post efficient mechanism would achieve (column "Second-best welfare" divided by column "First-best welfare").

Table 5.5: Efficiency of the TPAA mechanism for *uniformly* distributed QoS when the number of customers and the number of providers are increased equally.

$K = \min\{N, M\}$	**Second-best welfare**	**First-best welfare**	**Efficiency**
50	41.3419	65.9973	0.6264
100	86.6442	132.6719	0.6531
150	132.5587	199.3517	0.6649
200	178.4311	265.9973	0.6708
250	224.0339	332.6235	0.6735
300	270.0221	399.3313	0.6762
350	315.9203	466.0111	0.6779
400	361.5663	532.6608	0.6788
450	407.5797	599.3747	0.6800
500	453.1483	666.0127	0.6804
550	498.8646	732.6713	0.6809
600	544.6702	799.3493	0.6814
650	590.4833	866.0323	0.6818
700	636.1366	932.6922	0.6820
750	681.7975	999.3471	0.6822
800	727.6245	1066.0283	0.6826
850	773.0014	1132.6003	0.6825
900	818.9349	1199.3372	0.6828
950	865.0273	1266.0462	0.6833
1000	910.6372	1332.7073	0.6833

Figure 5.3 depicts the content of Table 5.5 graphically. It shows the mechanism's efficiency as the number of customer and the number of providers are increased equally; that is, $K = \min\{N, M\}$ is increased to infinity. For $N = M = 50$ agents (first data point), the efficiency has its lowest value at 0.6264. The curve then increases monotonically and approaches the asymptote at 0.6859 from below. This asymptotic efficiency of 0.6859 was derived analytically in formula (4.38).

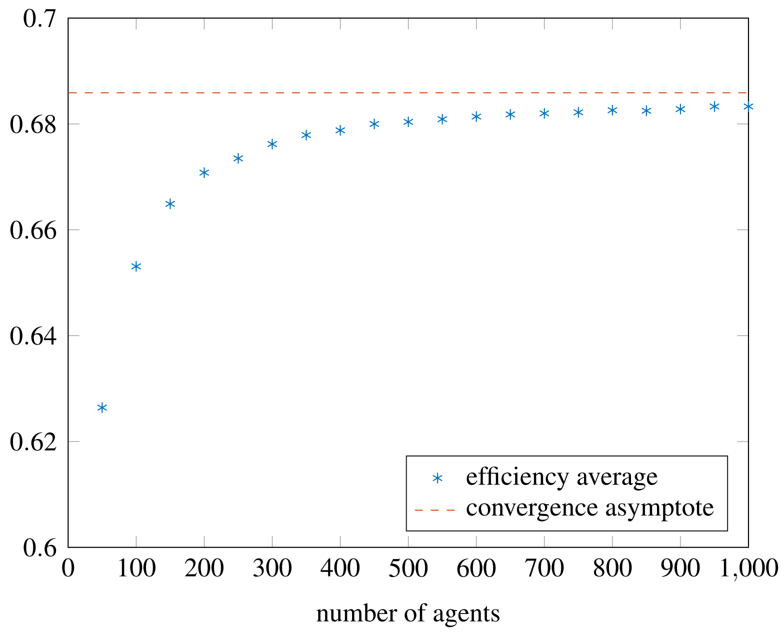

Figure 5.3: Asymptotic efficiency of the TPAA mechanism with symmetric agents and QoS realizations taken from the *uniform* distribution over the unit interval as $K \to \infty$.

5.2.2 Normally distributed, symmetric quality

This section provides the results obtained for a market, in which the QoS of customers and providers is normally distributed on the unit interval. Again, all agents are assumed to be symmetric.

5.2.2.1 Fixed number of providers

Table 5.6 shows the results obtained by the TPAA mechanism for a market with an increasing number of customers and the number of providers fixed to 10, 50, and 100. All agents are symmetric normal. That is, the QoS desired by customers and the QoS offered by providers are drawn from the normal distribution truncated to the unit interval with mean of $\mu = 0.5$ and standard deviation of $\sigma = 0.1$.

Table 5.6: Efficiency of the TPAA mechanism for *normally* distributed QoS with number of providers fixed to 10, 50, and 100.

$N = \lvert\mathscr{A}_C\rvert$	$M = 10$	$M = 50$	$M = 100$
2	0.6337	0.0782	0.0059
10	0.9253	0.5849	0.1737
20	0.9595	0.8190	0.5899
30	0.9612	0.8807	0.7556
40	0.9619	0.9096	0.8214
50	0.9628	0.9311	0.8574
60	0.9628	0.9593	0.8807
70	0.9631	0.9600	0.8973
80	0.9639	0.9605	0.9100
90	0.9640	0.9610	0.9205
100	0.9645	0.9613	0.9316

The content of Table 5.6 is illustrated graphically in Figure 5.4. The QoS of customer and providers are normally distributed. With 10 providers on the market, the lowest efficiency is 0.6337 for two customers. As the number of customers increases, the efficiency quickly increases and approaches the asymptote at 0.9735 for 10 or more customers. When 50 providers are on the market, the lowest efficiency is 0.0782 for two customers. The efficiency then increases and also approaches 0.9735 for 50 or more customers. When 100 providers offer their services to two customers, the efficiency is at its minimum of 0.0059. Like in the setting with 10 and 50 customers, the efficiency approaches 0.9735 for 100 customers (not shown in the graph). Similar to formula (5.3), the asymptotic efficiency can be evaluated analytically by

$$\frac{\int_a^b \Big(v(1,t) - c(1,t)\Big) f(t)\,dt}{\int_0^1 \Big(v(1,t) - c(1,t)\Big) f(t)\,dt} \approx 0.9735, \tag{5.4}$$

where $f(t) = \frac{1}{\sqrt{0.02\pi}} e^{-(t-0.5)^2/0.02}$ is the density of the normal distribution and $a = 0$ and $b = 0.6967$ are the solutions to $\psi_C(1,t) - \psi_P(1,t) = 0$.

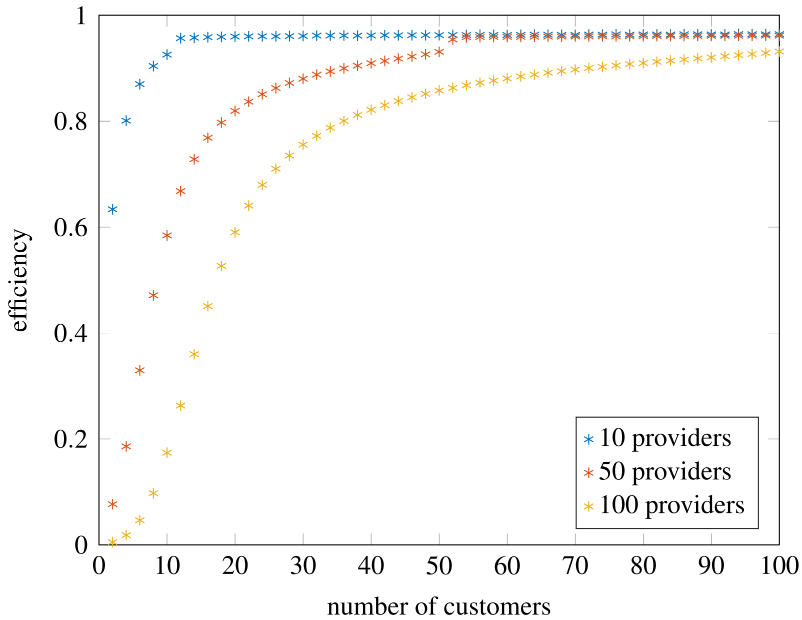

Figure 5.4: Efficiency of the TPAA mechanism with symmetric agents and QoS realizations taken from the *normal distribution* truncated to the unit interval with 10, 50, and 100 providers.

5.2.2.2 Increasing number of providers

Table 5.7 shows the efficiency of the TPAA mechanism when the number of customers and the number of providers are increased equally. All agents are symmetric normal; that is, all QoS realizations follow the normal distribution with mean of $\mu = 0.5$ and standard deviation of $\sigma = 0.1$.

Table 5.7: Efficiency of the TPAA mechanism for *normally* distributed QoS realizations when the number of customers and the number of providers are increased equally.

$K = \min\{N, M\}$	**Second-best welfare**	**First-best welfare**	**Efficiency**
10	13.2836	14.3503	0.9257
20	26.7292	28.7758	0.9289
30	40.1778	43.2020	0.9300
40	53.6373	57.6384	0.9306
50	67.0896	72.0721	0.9309
60	80.5403	86.5001	0.9311
70	94.0065	100.9298	0.9314
80	107.4848	115.3731	0.9316
90	120.9242	129.8062	0.9316
100	134.3688	144.2362	0.9316
110	147.7988	158.6618	0.9315
120	161.2170	173.0916	0.9314
130	174.7190	187.5380	0.9316
140	188.1320	201.9665	0.9315
150	201.5995	216.3959	0.9316
160	215.0549	230.8396	0.9316
170	228.5259	245.2781	0.9317
180	241.9595	259.7124	0.9316
190	255.3934	274.1296	0.9317
200	268.8460	288.5723	0.9316

The mechanism's efficiency as the number of customer and providers increases equally is shown in Figure 5.5. Here, the number of customers and the number of providers are increased at the same rate such that $M = N$ always holds. For 10 agents on each market side (first data point), the efficiency has its lowest value at 0.9257. The curve then increases monotonically and approaches the asymptote at 0.9317.

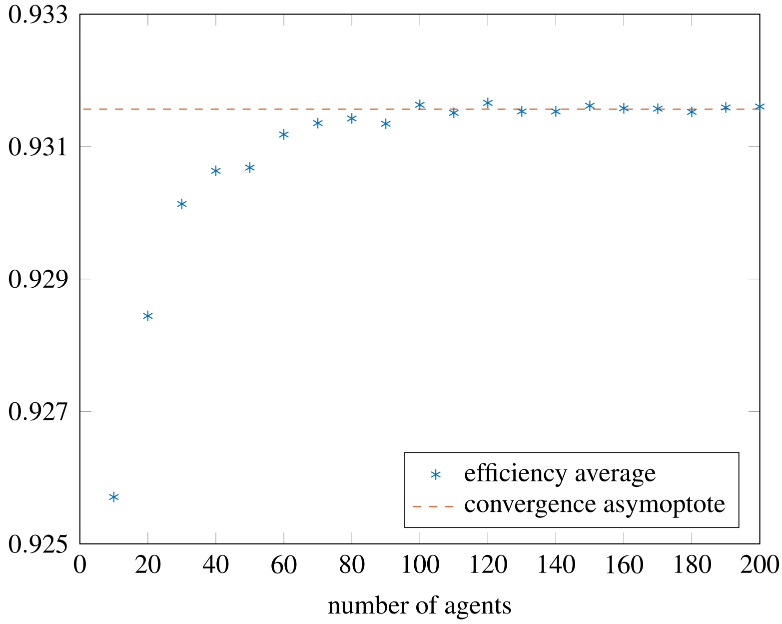

Figure 5.5: Asymptotic efficiency of the TPAA mechanism with symmetric agents and QoS realizations taken from the *normal* distribution truncated to the unit interval as $K \to \infty$.

Similar to formula (4.38), the asymptotic efficiency can be calculated as

$$\frac{\int_a^b \Big(v(t,t) - c(t,t)\Big) f(t)\, dt}{\int_0^1 \Big(v(t,t) - c(t,t)\Big) f(t)\, dt} \approx 0.9317, \tag{5.5}$$

where $f(t)$ is the density of the normal distribution, and $a = 0.3065$ and $b = 0.6692$ are the solutions to $\psi_C(t,t) - \psi_P(t,t) = 0$. Hence, in the limit market with normally distributed QoS, the TPAA mechanism excludes any customer-provider pair with QoS less than 0.3065 and greater than 0.6692. These exclusion lines are depicted in Figure 5.6. Obviously, for symmetric customers and providers with normally distributed QoS, the exclusion lines surround the mean at 0.5 almost evenly.

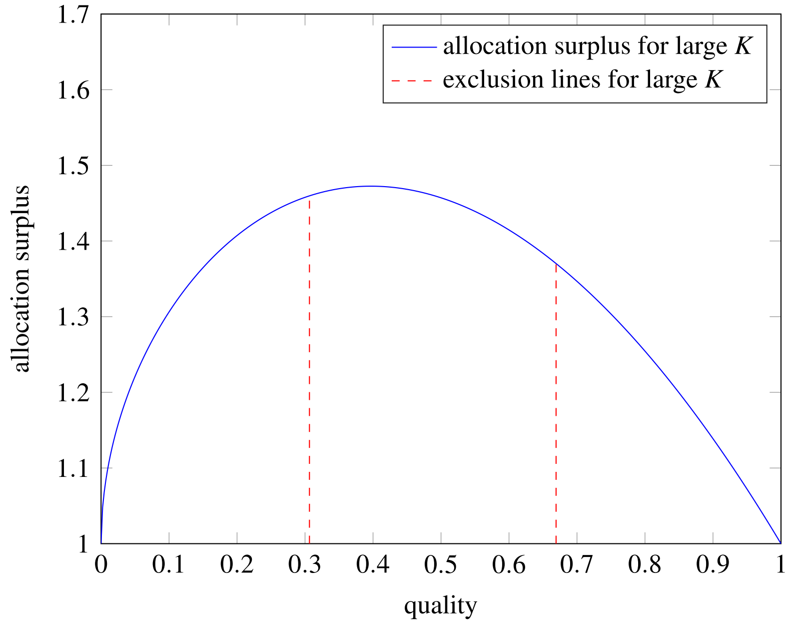

Figure 5.6: The solid curve shows the allocation surplus for $K \to \infty$, while only the QoS pairs between the dashed lines are matched under the normal distribution.

5.2.3 Beta-distributed, asymmetric quality

This section presents the results generated by the mechanism when demand and supply are asymmetrically beta-distributed. Like in the previous sections, the market involves both a fixed and an increasing number of providers.

5.2.3.1 Fixed number of providers

In this experiment, the QoS desired by customers is drawn from the beta distribution with $\alpha = 3$ and $\beta = 2$, while the QoS offered by providers is drawn from the beta distribution with $\alpha = 1$ and $\beta = 3$. Table 5.8 shows the results obtained by the TPAA mechanism for this asymmetric case.

Table 5.8: Efficiency of the TPAA mechanism for *beta-distributed* QoS with number of providers fixed to 10, 50, and 100.

$N = \|\mathscr{A}_C\|$	$M = 10$	$M = 50$	$M = 100$
2	0.5106	0.0332	0.0012
10	0.8982	0.3842	0.0438
20	0.9352	0.7575	0.3461
30	0.9337	0.8513	0.6476
40	0.9317	0.8884	0.7681
50	0.9301	0.9083	0.8233
60	0.9295	0.9429	0.8549
70	0.9283	0.9419	0.8757
80	0.9279	0.9414	0.8902
90	0.9273	0.9403	0.9013
100	0.9267	0.9397	0.9097

The graphical representation of the content in Table 5.8 is shown in Figure 5.7. The lowest efficiency of 0.5106 is obtained when 10 providers offer their services to 2 customers. Once the number of customers increases to 20, the efficiency increases drastically up to a value of 0.9352. For 30 and more customers, however, the efficiency of the mechanism decreases at a very slow rate, reaching about 0.9267 when 100 customers demand services. A similar phenomenon can be observed for a market with 50 providers. Starting with a very low efficiency of 0.0332, the maximal efficiency of 0.9429 is obtained when 60 customers are in the market. Then, a small, constant decrease in efficiency follows until 0.9397 is reached for 100 customers. If 100 providers offer their services, the efficiency is monotone for the considered range of customers.

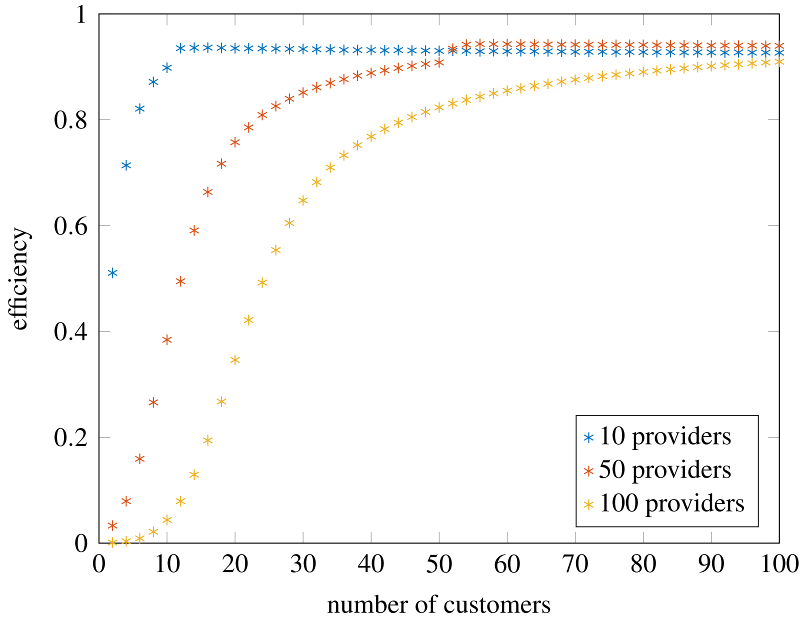

Figure 5.7: Efficiency of the TPAA mechanism with asymmetric agents and QoS realizations taken from the *beta distribution* with 10, 50, and 100 providers.

For markets with a fixed number of providers and many customers, the asymptotic efficiency is computed analytically as follows. If a large number customers desire services from a fixed number of providers, the QoS desired by customers converges to 1. Therefore, the efficiency is given by

$$\frac{\int_a^b \Big(v(1,t) - c(1,t)\Big) f(t)\, dt}{\int_0^1 \Big(v(1,t) - c(1,t)\Big) f(t)\, dt} \approx 0.9111, \tag{5.6}$$

where $f(t) = 3(t-1)^2$ is the density of the beta distribution for providers with $\alpha = 1$ and $\beta = 3$. Similar to the previous experiments, $a = 0$ and $b = 0.6376$ are the solutions to $\psi_C(1,t) - \psi_P(1,t) = 0$. Therefore, in comparison to the efficiency values shown in Table 5.8, the asymptote at 0.9111 is approached from above. Additional experiments with larger numbers of customers confirm this phenomenon.

5.2.3.2 Increasing number of providers

Table 5.9 shows the efficiency of the TPAA mechanism when the number of customers and the number of providers are increased equally. In this case, customers are providers are asymmetric. That is, customers desire their QoS according to the beta distribution with $\alpha = 3$ and $\beta = 2$. Hence, customers are more likely to desire high quality services. In contrast, the QoS offered by providers is drawn from the beta distribution with $\alpha = 1$ and $\beta = 3$. Thus, providers are likely to offer low offered quality because they anticipate lower provision cost (cf. Section 5.1).

Table 5.9: Efficiency of the TPAA mechanism for *beta-distributed* QoS realizations, when the number of customers and the number of providers are increased equally.

$K = \min\{N, M\}$	**Second-best welfare**	**First-best welfare**	**Efficiency**
10	12.3967	13.8118	0.8975
20	25.2212	27.8721	0.9049
30	38.0437	41.9405	0.9071
40	50.8472	56.0192	0.9077
50	63.6848	70.0978	0.9085
60	76.5118	84.1616	0.9091
70	89.3304	98.2379	0.9093
80	102.1583	112.3244	0.9095
90	114.9787	126.3858	0.9097
100	127.7763	140.4655	0.9097
110	140.6281	154.5491	0.9099
120	153.4471	168.6043	0.9101
130	166.2508	182.6889	0.9100
140	179.1028	196.7532	0.9103
150	191.9128	210.8461	0.9102
160	204.7530	224.9219	0.9103
170	217.5585	238.9911	0.9103
180	230.3999	253.0731	0.9104
190	243.2257	267.1530	0.9104
200	256.0136	281.2237	0.9104

The contents of Table 5.9 are depicted graphically in Figure 5.8. When the number of customers and the number of providers are increased at the same rate, the mechanism's efficiency is monotone throughout the complete range of QoS values.

For $N = M = 10$ agents on each market side (first data point), the efficiency has its lowest value at 0.8975. The curve then increases monotonically and reaches its maximum at 0.9104 for 200 agents on either market side.

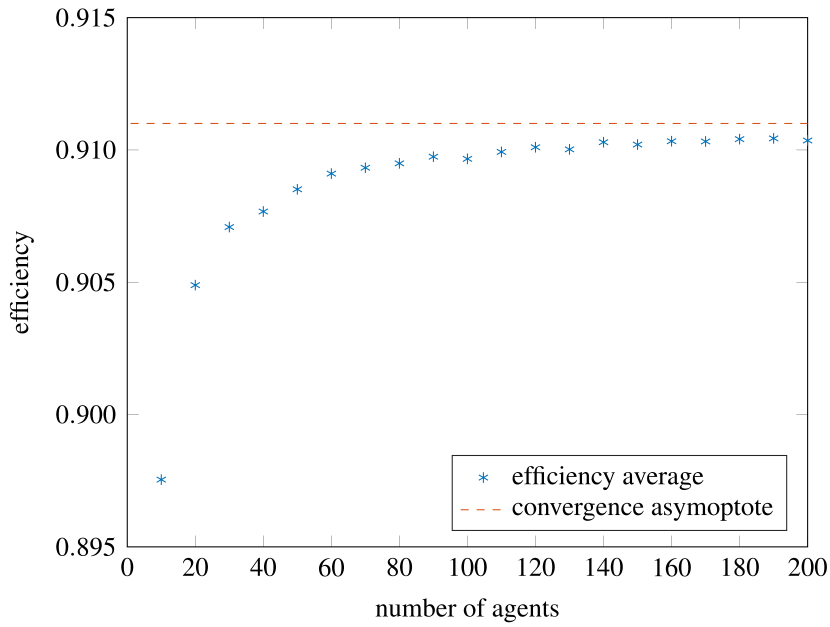

Figure 5.8: Asymptotic efficiency of the TPAA mechanism with asymmetric agents and QoS realizations taken from the *beta* distribution as $K \to \infty$.

To obtain an analytic expression for the asymptotic efficiency, the asymmetry of customers and providers in this experiment must be considered. For this purpose, it is necessary to calculate the allocation surplus in the limit market as $K \to \infty$. As discussed in Section 4.2.3, the TPAA mechanism converges to deterministic quantile matching. In the limit as $K \to \infty$, the empirical quantile function is asymptotically equivalent to the k-th order statistic. Therefore, the expected QoS offered by providers is a function of the two cumulative distributions F and H defined in Section 3.2.1. That is, for any provider's QoS in the limit market, $\sigma \to H^{-1}(F(\theta))$. Because the QoS desired by customers is beta-distributed with $\alpha = 3$ and $\beta = 2$, their cumulative distribution is $F(\theta) = (4-3\theta)\theta^3$. Similarly, the QoS offered by providers is $H(\sigma) = (\sigma^2 - 3\sigma + 3)\sigma$. Hence, inverting H and substituting $\sigma = F(\theta)$ yields

$$H^{-1}(F(\theta)) = ((4-3\theta)\theta^3 - 1)^{\frac{1}{3}}. \tag{5.7}$$

Using this expression, the allocation surplus for the limit market is given by

$$\begin{aligned} v(\theta, H^{-1}(F(\theta))) - c(\theta, H^{-1}(F(\theta))) \\ &= \sqrt{\theta} - 2\Big(((4-3\theta)\theta^3 - 1)^{\frac{1}{3}} - x + 1\Big)^2 \\ &\quad - \Big(((4-3\theta)\theta^3 - 1)^{\frac{1}{3}} + 1\Big)^2 + 1. \end{aligned} \tag{5.8}$$

Figure 5.9 depicts the allocation surplus displayed in equation (5.8) for large markets as a function of the QoS θ desired by customers. The QoS of the associated provider in that match is then given by $\sigma = H^{-1}(F(\theta))$. The shape of the curve differs from the shapes observed for the uniform distribution (cf. Figure 4.4) and the normal distribution (cf. Figure 5.6) in the number of local maxima. Unlike the latter two graphs, the graph obtained for beta-distributed QoS shows two local maxima. The first maximum emerges at a customer QoS of $\theta = 0.3663$ with provider QoS of $\sigma = 0.05$. The second maximum appears at $\theta = 0.8107$ with $\sigma = 0.4521$. Thus, the highest allocation surplus is generated when a customer with demand 0.8107 and a provider with offer $\sigma = 0.4521$ are matched together in the limit market. Certainly, the existence of two maxima reflects the asymmetry of customers and providers in this setting (cf. discussion in Section 5.3.4).

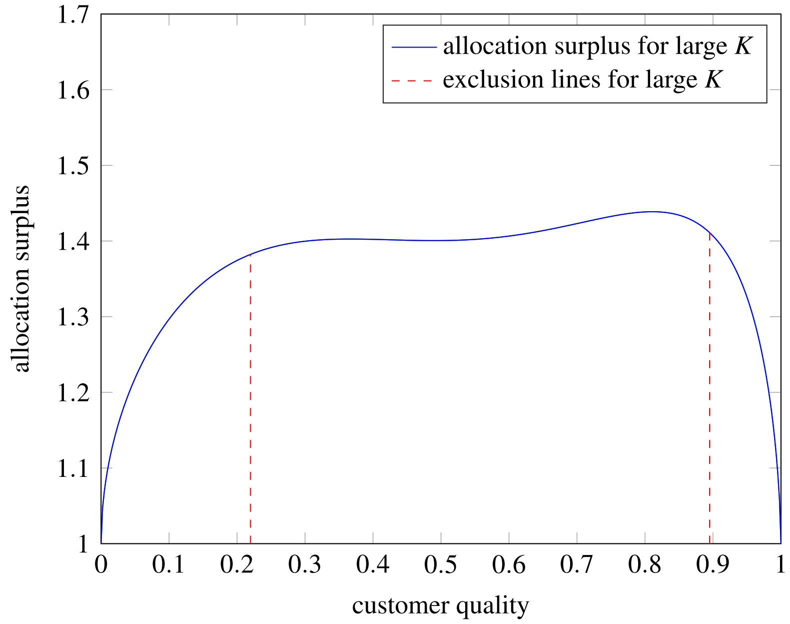

Figure 5.9: The solid curve shows the allocation surplus for $K \to \infty$ as a function of customer QoS, while only the desired QoS between the dashed lines is matched under the beta distribution.

Similar to the experiments with uniform and normal QoS, the asymptotic efficiency of the mechanism for beta-distributed QoS can be calculated analytically. The ratio between the first-best and the second-best outcome is given by

$$\frac{\displaystyle\int_a^b \Big(v(t,H^{-1}(F(t))) - c(t,H^{-1}(F(t)))\Big) f(t)\,dt}{\displaystyle\int_0^1 \Big(v(t,H^{-1}(F(t))) - c(t,H^{-1}(F(t)))\Big) f(t)\,dt} \approx 0.9110, \tag{5.9}$$

where $f(t) = 12(1-t)t^2$ is the density function of the beta distribution for $\alpha = 3$ and $\beta = 2$. The bounds of integration $a = 0.2198$ and $b = 0.8950$ are the solutions to the exclusion condition $\psi_C(t,H^{-1}(F(t))) - \psi_P(t,H^{-1}(F(t))) = 0$. The efficiency of 0.9110 derived in (5.9) is shown as the convergence asymptote in Figure 5.8.

5.3 Discussion

This section discusses the insights that can be obtained from the results presented in Section 5.2. Particular attention is paid to the fact that the asymptotic efficiency of the TPAA mechanism is bounded away from 100% for each experiment. For this purpose, the economic concept of each agent's informational smallness is discussed in detail. In conjunction with the informational smallness of the agents, the size of each agent's informational rent is of particular interest. Subsequently, this section sheds light on the non-vanishing difference between a customer's valuation and a provider's cost that is necessary for successful matching. Moreover, a number of observations on the results obtained for the different probability distributions are emphasized. Finally, this section discusses possible implementations of the TPAA mechanism in dominant strategies.

5.3.1 Informational smallness of customers and providers

The results obtained in all experiments imply that although the mechanism's asymptotic efficiency can be quite high, it is bounded away from 100%. This finding differs significantly from the results described by other work in mechanism design theory. Single-unit double auctions, for instance, achieve a socially optimal outcome in the limit market under some assumptions (Wilson, 1985). Similar results hold for multilateral trading mechanisms. Unlike the bilateral trading mechanism proposed by Myerson and Satterthwaite (1983), which is not ex post optimal, the multilateral mechanism studied by Gresik and Satterthwaite (1989) arrives at an asymptotic efficiency of 100% for limit markets. They show that the market's relative inefficiency is at most of the order $\log K/K^2$ for large K, where the number of customers and the number of providers are increased at equal rates. Therefore, as $K \to \infty$, any inefficiencies disappear.

The mechanism studied in this research, however, indicates that the presence of many customers and providers does *not* eliminate the mechanism's inefficiency caused by the asymmetry of information. The reason for this phenomenon is related to the *informational smallness* of the agents. Loosely speaking, an agent is informationally small if the incremental impact of that agent's private information (i.e., desired/offered QoS) on the demand of every electronic service is "small", given the information of other agents. The concept of informational smallness has been studied by Gul and Postlewaite (1992) in the context of replicated economies. McLean and Postlewaite (2002, 2004) extend Gul and Postlewaite's model to a more general framework by providing a precise formalization of an agent's informational size.

In principle, the asymptotic inefficiency observed for the TPAA mechanism could disappear if certain conditions were satisfied. Gul and Postlewaite (1992) identify such conditions sufficient for eliminating inefficiencies in replicated economies. However, for certain settings with asymmetric information, these conditions are invalid. Therefore, the associated mechanisms cannot achieve full efficiency even when the number of agents is large. One such example is the replication of the classic bilateral trading mechanism studied by Myerson and Satterthwaite (1983), in which only a specific match produces pairwise private surpluses. The two-trader case illustrates the impossibility of designing incentive compatible mechanisms that are ex post efficient. In fact, the associated second-best mechanism reaches an efficiency of 84.36% for uniformly distributed types (Gresik and Satterthwaite, 1983). If this bilateral case were replicated with independent valuations across pairs such that only specific matches generate mutual value, the inefficiency due to asymmetric information might not vanish (Gul and Postlewaite, 1992). The presence of additional pairs in the market would not impact the problem faced by any particular pair for allocating the object between the two traders. This example demonstrates that even in the presence of many agents (and consequently many objects), the asymptotic efficiency of the mechanism may still be bounded away from 100%.

The proposed model of the TPAA mechanism can be regarded as a specific replication of the Myerson-Satterthwaite bilateral trading model. Surplus is produced only if a specific pair of two agents is matched together. Only this particular agent pair cares about the electronic service to be allocated. Adding more agent pairs to this economy does not change the incremental impact of each agent's private information on the demand of the electronic service. Thus, the agents in the proposed model are *not* informationally small even in large markets. In the model proposed in this thesis, each provider offers a service of differentiated QoS, and each customer desires a service of differentiated QoS. This feature prevent the agents from becoming informationally small as the market becomes large. To this end, the absence of informational smallness is the reason why the asymptotic inefficiency due to the asymmetry of information does not disappear in the TPAA mechanism.

The numerical example discussed in Section 4.2.3 shows that the efficiency of the TPAA mechanism for large markets is 68.59%, given uniformly distributed QoS on the unit interval. This efficiency, however, is significantly lower than that of the bilateral trading mechanism (84.36% as reported by Gresik and Satterthwaite (1983)). If the latter economy were to be replicated, the number of commodities would approach infinity, while the agents' valuations would still remain independent. This research, however, focuses on a framework in which an agent's utility depends upon the types of all other agents. In particular, a customer's utility is maximized when his desired QoS matches the offered QoS delivered by a provider. Similarly, providers maximize their utility by delivering exactly the QoS desired by customers. This interdependence of the agents' utilities causes the additional efficiency loss of over 15.77% as compared to the outcome found by Myerson and Satterthwaite (1983). This observation suggests that agents with interdependent utilities have greater incentives to misreport their QoS in order to win a better allocation than agents in the (replicated) bilateral trading mechanism with independent valuations.

5.3.2 Size of informational rents

Although the structure of the model proposed in this work differs from that of McLean and Postlewaite (2002), their notion of an agent's informational size can, to some extent, still be applied to this model. In fact, the *informational rent* of an agent quantifies the amount of compensation that must be paid to that agent to induce honest reporting of their QoS. The agents' informational rents are linked to their informational size (McLean and Postlewaite, 2004). In this model, the agents' informational rents are given by the second summand in the transfer functions (4.24) and (4.25). For the case of asymmetric, beta-distributed QoS (cf. Experiment 3), the informational rent $S_C^\infty(\theta)$ of customers in the limit market is given by

$$\begin{aligned} S_C^\infty(\theta) &= \int_0^\theta \frac{\partial}{\partial t} v(t, H^{-1}(F(t)))dt \\ &= 2\theta - \theta^2 + \sqrt{\theta} - \Big(\theta^3(3\theta-4)+1\Big)^{\frac{1}{3}} \Big(2\theta + (\theta^3(3\theta-4)+1)^{\frac{1}{3}} - 2\Big) - 1. \end{aligned} \tag{5.10}$$

Similarly, the informational rent $S_P^\infty(\sigma)$ of providers is

$$\begin{aligned} S_P^\infty(\sigma) &= \int_\sigma^1 \frac{\partial}{\partial t} c(t, H^{-1}(F(t)))dt \\ &= 2\sigma - 2\sigma^2 - \Big(\sigma^3(3\sigma-4)+1\Big)^{\frac{1}{3}} \Big(2\sigma + (\sigma^3(3\sigma-4)+1)^{\frac{1}{3}} - 2\Big). \end{aligned} \tag{5.11}$$

Hence, both quantities are bounded away from zero when the number of agents becomes large on both market sides. Figure 5.10 depicts the informational rents of customers and providers as functions of the QoS. This example with beta-distributed, asymmetric QoS on $[0,1]$ shows that the amount of compensation for honest reporting does not disappear even when the market size goes to infinity. In this sense, the informational size of the agents does not not vanish for large markets.

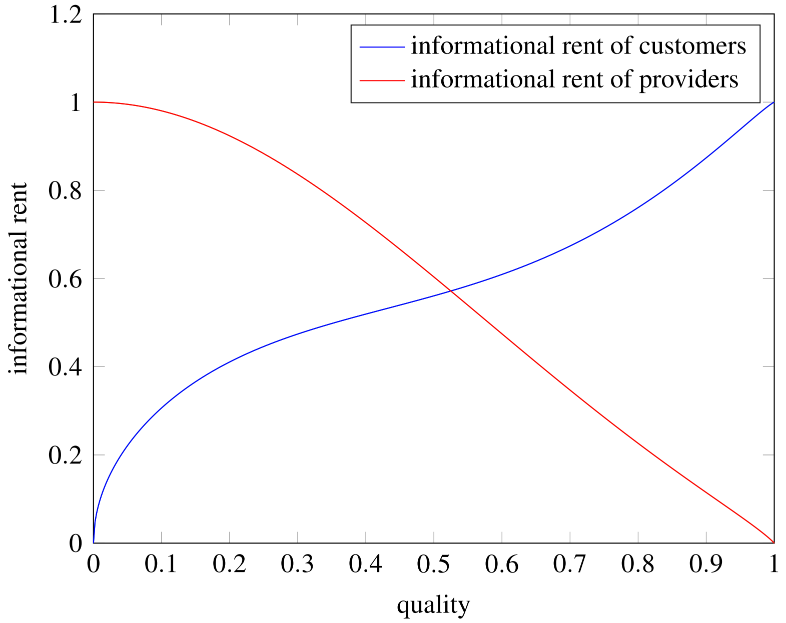

Figure 5.10: Informational rents of asymmetric customers and providers with beta-distributed QoS in the limit market.

5.3.3 Non-vanishing difference in valuation and cost

As discussed in Section 4.2.1, the necessary condition for any two agents to match in the limit market is $\psi(\theta,\theta) \geq 0$, when the agents use symmetric, uniformly distributed QoS. This difference in virtual reservation values required by the allocation rule (4.23) does not vanish even when the number of traders is large. Based on the parameters of the example in Section 4.2.3, inequality $\psi(\theta,\theta) \geq 0$ can be written as

$$v(\theta,\theta) - c(\theta,\theta) \geq 2\theta^2 + \frac{1-\theta}{2\sqrt{\theta}} > 0.7915. \tag{5.12}$$

This difference in valuation and provision cost must be satisfied to warrant service allocation in large markets. In particular, the right side of (5.12) is bounded away from zero by 0.7915 for all $\theta \in (0,1]$ even when $K \to \infty$. In the asymptotically efficient mechanism proposed by Gresik and Satterthwaite (1983), the difference in reservation values vanishes at the same rate as $1/2K$ approaches zero. Again, the presence of this non-vanishing threshold (5.12) confirms that the TPAA mechanism cannot attain full efficiency even when the number of traders is large.

5.3.4 Efficiency properties under different distributions

The mechanism's efficiency properties arising in the three probability distributions uniform, normal, and beta show a number of considerable distinctions. While the asymptotic efficiency with symmetric, uniformly distributed QoS remains below 70% for equal number of agents on both market sides, the efficiency for symmetric normal and asymmetric beta-distributed QoS exceeds 90%. The following paragraphs discuss the issues that emanate from these observations.

5.3.4.1 Uniform distribution

The results obtained for symmetric agents with uniformly distributed QoS on the unit interval show that the efficiency of the mechanism is non-monotonic when the number of providers is fixed. However, these non-monotonicities disappear once the number of customers becomes large because the efficiency converges to 0.8397. In particular, a market with 60 customers and 10 providers achieves the highest efficiency, followed by a constant decrease towards the convergence asymptote. The economic intuition behind this result is not sufficiently clear and requires further research. Interestingly, Muthoo and Mutuswami (2005) report on similar non-monotonicities that occur in settings with discrete quality types. An economic explanation for this phenomenon, however, is not provided.

In contrast to fixed provider numbers, there are no monotonicity issues when the number of customers and providers are increased equally. In fact, the efficiency values smoothly approach the asymptote from below. However, the asymptotic efficiency remains quite low at 68.59%. With QoS uniformly scattered across the unit interval, there is more leeway for customers and providers to distort their true type. Hence, it is comparatively expensive for the mechanism to implement incentive compatibility.

5.3.4.2 Normal distribution

The results obtained in the setting with symmetric customers and providers and normally distributed QoS (mean of 0.5 and standard deviation of 0.1) are different from the uniform case. As in the uniform case, the asymptotic efficiency is also bounded away from 100%. However, it is significantly higher than that for the uniform distribution. In the considered range of up to 100 customers, non-monotonicities in efficiency do not arise for fixed numbers of providers. Yet this observation may not apply for a higher number of customers. It is assumed that slight non-monotonicities may also occur with normally distributed qualities. Further investigation concerning these non-monotonicities, however, is out of the scope of this thesis and could be the subject of future research.

For an equal number of agents on either market side, normally distributed QoS entails an asymptotic efficiency of 93.15% as $K \to \infty$. When the QoS desired/offered by the agents is drawn from the normal distribution, the associated realizations are scattered around its mean of 0.5 with exclusion lines beyond the standard deviation of ± 0.1 (cf. Figure 5.6). In such situations, there is more mass in the center, which means that the agents' incentive to distort their QoS for receiving a better allocation is smaller. Figure 5.11 depicts this distortion due to the presence of private information for symmetric agents whose QoS is drawn from the normal distribution.

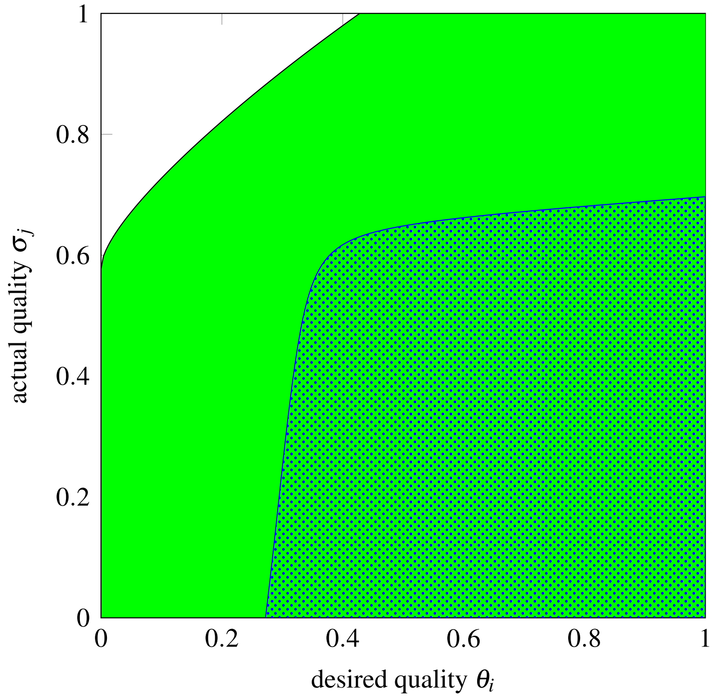

Figure 5.11: The shaded area (solid green) contains all eligible, *normally distributed* QoS pairs that warrant service allocation in the first-best mechanism. The dotted area contains the corresponding quality pairs in the second-best mechanism.

5.3.4.3 Beta distribution

Unlike Experiments 1 and 2, the third experiment considers *asymmetric* customers and providers. This setting illustrates a market, in which customers tend to desire high-quality QoS, and providers are more likely to offer low quality in order to reduce provision costs. The associated relative probabilities are shown in Figure 5.1. Because the agents are asymmetric in their QoS, the boundary for successful matching in the TPAA mechanism depicted in Figure 5.12 is not symmetric. For this asymmetric, beta-distributed market, the expected QoS desired by customers is 0.6, and the expected QoS offered by providers is 0.25. The same asymmetry appears in Figure 5.9, where the allocation surplus in the limit market has two local maxima.

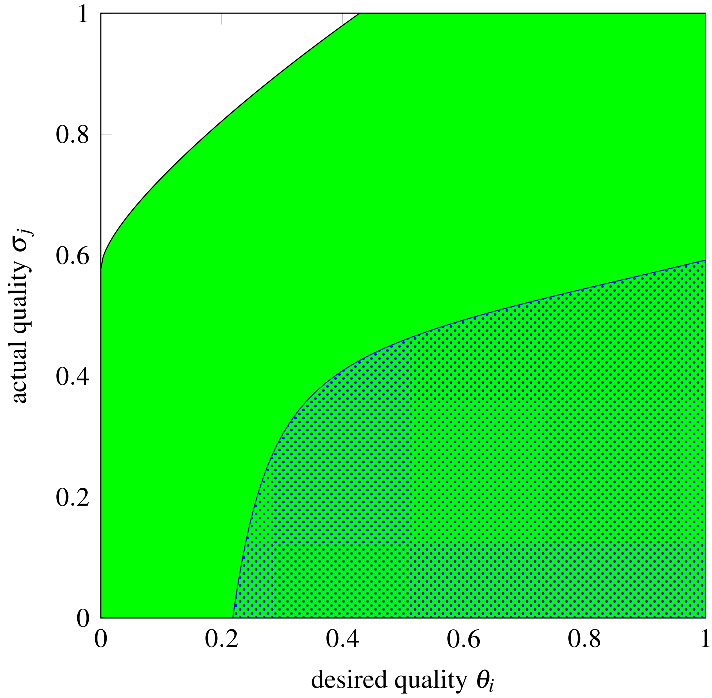

Figure 5.12: The shaded area (solid green) contains all eligible, *beta-distributed* QoS pairs that warrant service allocation in the first-best mechanism. The dotted area contains the corresponding quality pairs in the second-best mechanism.

Similar to the case with uniformly distributed QoS, the efficiency of the mechanism exhibits non-monotonic behavior when the number of providers is fixed. If 10 providers are in the market, the decrease in efficiency starts already with 30 customers (as opposed to 60 customers in the uniform case) and decreases by almost 1%. These differences may result from the asymmetry of customers and providers in the beta-distributed case.

Although customers and providers are asymmetric with different expected QoS, the asymptotic efficiency of the mechanism is still quite high at 91.10%. For *Beta*$(3,2)$, the QoS desired by customers is concentrated around $\theta = 2/3$, which is the maximum of the density function (cf. Section 5.1.1). To this end, a customer is faced with higher competition because the QoS realizations of other customers are clustered around $2/3$. Therefore, customers are less incentivized to distort their true desired QoS. On the supply side, a provider with decreasing density according to *Beta*$(1,3)$ competes with other providers to offer low-quality services. Because the mechanism lowers the probability that high-quality providers will be allocated, a provider's incentive to distort his offered QoS is small (cf. Section 4.2.3). As a result, both customers and providers have little incentive to misrepresent their private quality information. Thus, the mechanism's efficiency is relatively high.

5.3.5 Implementation in dominant strategies

The proposed mechanism can also be implemented in dominant strategies with the following considerations. In general, imposing incentive compatibility constraints for dominant-strategy mechanisms instead of Bayesian mechanisms considerably restricts the set of implementable mechanisms. However, Mookherjee and Reichelstein (1992) identify mechanism design problems for which Bayesian incentive compatibility can be replaced by the more stringent requirement of dominant strategy incentive compatibility without any losses. Since the allocation rule defined in Theorem 4.1 satisfies the monotonicity conditions of Mookherjee and Reichelstein (1992), there is no loss from replacing Bayesian equilibrium constraints with dominant strategy requirements. Hence, the TPAA mechanism proposed in this research is implementable in dominant strategies.

Chapter 6

Conclusions

Now all has been heard; here is the conclusion of the matter: Fear God and keep His commandments, for this is the duty of all mankind.

Ecclesiastes 12:13, The Bible
(New International Version)

This chapter highlights the contributions and key findings obtained in this work and outlines opportunities for future research.

6.1 Contributions

This research addresses the problem of service allocation under double-sided competition with private quality information. Using a mechanism design framework, this work presents a mechanism for allocating electronic services with privately known QoS and an analysis of its asymptotic efficiency properties. This thesis makes the following two specific contributions to research in mechanism design theory.

First, this research adds to the literature by developing a double-sided mechanism for the allocation of electronic services with private quality information from a social welfare perspective. Prior research has mainly focused on the maximization of the intermediary's profit (Che, 1993; Bichler and Kalagnanam, 2005; Johnson, 2013), whereas social welfare maximization subject to budget balance has heretofore received little attention. Because marketplaces for electronic services must facilitate distributed decision-making among the agents, budget-balanced mechanisms should be investigated. Therefore, this research extends prior research that has solely focused on the optimal allocation rules from the perspective of a profit-maximizing intermediary. This extension includes the design of the optimal mechanism from the perspective of a social planner who attempts to maximize the expected social welfare generated by the agents. Because private quality information is present in the analysis, this optimal mechanism is a second-best mechanism for allocating electronic services. Specifically, the allocation mechanism is incentive compatible, individually rational, and balances the budget.

Second, research on the efficiency properties of second-best allocation mechanisms is still lacking and has only considered either identical objects (Gresik and Satterthwaite, 1989; Rustichini et al., 1994) or discrete quality levels without deriving the optimal allocation rules and payments (Muthoo and Mutuswami, 2005, 2011). This research fills this gap in the literature by studying the efficiency properties of the second-best mechanism in three distinct simulation experiments. The first two experiments consider a market with symmetric customers and providers. In these experiments, the QoS of the agents is drawn from the uniform and the normal distribution. The third experiment assumes asymmetric agents with beta-distributed QoS. Throughout all experiments, this research finds that the asymptotic efficiency of the mechanism is bounded away from 100% even when number of agents becomes large. This finding is related to the economic concept of informational smallness, which is defined as the incremental impact of an agent's QoS on the demand of an electronic service. In the proposed model, each provider offers a service of distinct QoS, and each customer demands a service of distinct QoS. It is this feature of differentiated service quality that prevents the agents from becoming informationally small as the market becomes large.

If each agent's private information about QoS follows the uniform distribution, the mechanism must tolerate an efficiency loss of more than 31% for an increasing number of customers and providers. In contrast, if private quality information is normally distributed among agents, this research finds that the mechanism's asymptotic inefficiency due to the asymmetry of information can be reduced to about 7% as the market size increases on both sides. With asymmetric, beta-distributed QoS, the mechanism arrives at an asymptotic efficiency of more than 91%. These findings are crucial to social planners because when designing service allocation with double-sided competition, they can obtain an accurate estimation of potential efficiency losses that arise from asymmetric information about QoS. Conversely, the social planner can ensure that every allocation decision is made by the agents only. Hence, the emerging distributed mechanism implementation circumvents the need for an external, independent decision maker.

6.2 Future work

Future research might be pursued in two directions. First, in the current model, QoS is restricted to a one-dimensional parameter. This limitation might be insufficient for designing electronic service allocation with multiple QoS attributes. Thus, the model could be extended to multi-attribute services by integrating scoring rules for double-sided markets. This change would affect the optimality results and require revisiting the model's properties. Second, the model's efficacy could be studied on a greater scale to provide a more comprehensive quantitative evaluation. For this purpose, the mechanism could be studied in simulation experiments with actual or synthetic data that more realistically reflect markets for electronic services including QoS and interdependent utility functions.

Bibliography

Abramowitz, M. and Stegun, I. A., editors. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover Publications, 1972.

Akerlof, G. The market for lemons: Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3):488–500, 1970.

Archer, A., Papadimitriou, C., Talwar, K., and Tardos, E. An approximate truthful mechanism for combinatorial auctions with single parameter agents. *Internet Mathematics*, 1(2):129–150, 2005.

Armbrust, B., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., Lee, G., Patterson, D., and Rabkin, A. A view of cloud computing. *Communications of the ACM*, 53: 50–58, 2010.

Armstrong, M. Competition in two-sided markets. *RAND Journal of Economics*, 37 (3):668–691, 2006.

Asker, J. and Cantillon, E. Properties of scoring auctions. *RAND Journal of Economics*, 39(1):69–85, 2008.

Avnet Cloud Marketplace. The Avnet Cloud Marketplace, 2016. URL `http://www.avnetpromotions.com/cloudmarketplace/`. [Online; accessed 12-February-2016].

Bagnoli, M. and Bergstrom, T. Log-concave probability and its applications. Working Papers 89-23, University of Michigan, 1989.

Bazerman, M. and Gillespie, J. Betting on the future: The virtues of contingent contracts. *Harvard Business Review*, 77(5):155–160, 1999.

Bhargava, H. K. and Sun, D. Pricing under quality of service uncertainty: Market segmentation via statistical QoS guarantees. *European Journal of Operational Research*, 191:1189–1204, 2008.

Bichler, M. and Kalagnanam, J. Configurable offers and winner determination in multi-attribute auctions. *European Journal of Operational Research*, 160(2):380–394, 2005.

Bichler, M., Lee, J., Lee, H., and Kalagnanam, J. Resource allocation algorithms for electronic auctions: A framework design. In *Proceedings of the 3rd International Conference on Electronic Commerce and Web Technologies (EC-Web)*, Aix-en-Provence, France, 2002.

Blau, B., van Dinther, C., Conte, T., Xu, Y., and Weinhardt, C. How to coordinate value generation in service networks? a mechanism design approach. *Business and Information Systems Engineering*, 1(5):343–356, 2009.

Blau, B., Conte, T., and van Dinther, C. A multidimensional procurement auction for trading composite services. *Electronic Commerce Research and Applications*, 9:460–472, 2010.

Bockstedt, J. and Goh, K. Seller strategies for differentiation in highly competitive online auction markets. *Journal of Management Information Systems*, 28(3):235–267, 2011.

Börgers, T. *An Introduction to the Theory of Mechanism Design*. Oxford University Press, 2015.

Branco, F. The design of multidimensional auctions. *RAND Journal of Economics*, 28(1):63–81, 1997.

Buyya, R., Yeo, C. S., Venugopal, S., Broberg, J., and Brandic, I. Cloud computing and emerging it platforms: Vision, hype, and reality for delivering computing as the 5th utility. *Future Generation Computer Systems*, 25(6):599–616, 2009.

Chambers, C. P. and Echenique, F. Supermodularity and preferences. *Journal of Economic Theory*, 144(3):1004–1014, 2009.

Chang, Y. B. and Gurbaxani, V. Information technology outsourcing, knowledge transfer, and firm productivity: An empirical analysis. *MIS Quarterly*, 36(4): 1043–1063, 2012.

Che, Y.-K. Design competition through multidimensional auctions. *RAND Journal of Economics*, 24:668–680, 1993.

Chevaleyre, Y., Dunne, P. E., Endriss, U., Lang, J., Lemaitre, M., Maudet, N., Padget, J., Phelps, S., Rodriguez-Aguilar, J. A., and Sousa, P. Issues in multiagent resource allocation. *Informatica*, 30:3–31, 2006.

Cripps, M. and Swinkels, J. Efficiency of large double auctions. *Econometrica*, 74 (1):47–92, 2006.

Dasgupta, P., Hammond, P., and Maskin, E. The implementation of social choice rules: Some general results on incentive compatibility. *Review of Economic Studies*, 46:185–216, 1979.

Dasgupta, S. and Spulber, D. F. Managing procurement auctions. *Information Economics and Policy*, 4(1):5–29, 1990.

David, H. A. and Nagaraja, H. N. *Order Statistics*. Wiley Series in Probability and Statistics, 2003.

Denoeud-Belgacem, L., Gourdin, E., Krishnaswamy, R., and Ouorou, A. Combinatorial auctions for exchanging resources over a grid network. *Annals of Telecommunications*, 65(11):689–704, 2010.

Egri, P. and Váncza, J. A distributed coordination mechanism for supply networks with asymmetric information. *European Journal of Operational Research*, 226 (3):452–460, 2013.

Eisenmann, T. R., Parker, G., and Alstyne, M. W. V. Strategies for two sided markets. *Harvard Business Review*, 84(10):92–101, 2006.

Engel, Y., Wellman, M. P., and Lochner, K. M. Bid expressiveness and clearing algorithms in multiattribute double auctions. In *Proceedings of the 7th ACM conference on Electronic commerce*, pages 110–119, 2006.

Fudenberg, D. and Tirole, J. *Game Theory*. MIT Press, Cambridge, MA., 1993.

Gadrey, J. The characterization of goods and services: An alternative approach. *Review of Income and Wealth*, 46(3):369–387, 2000.

Gartner. *Forecast Analysis: IT Outsourcing, Worldwide, 2Q13 Update*. www.gartner.com, Stamford, USA, 2013.

Gibbard, A. F. Manipulation of voting schemes: A general result. *Econometrica*, 41: 587–602, 1973.

Greenberg, A., Hamilton, J., Maltz, D. A., and Patel, P. The cost of a cloud: research problems in data center networks. *ACM SIGCOMM Computer Communication Review*, 39(1):68–73, 2009.

Gresik, T. A. and Satterthwaite, M. A. The number of traders required to make a market competitive: the beginnings of a theory. *CMSEMS DP.*, 551:537–557, 1983.

Gresik, T. A. and Satterthwaite, M. A. The rate at which a simple market converges to efficiency as the number of traders increases: An asymptotic result for optimal trading mechanisms. *Journal of Economic Theory*, 48(1):304–332, 1989.

Gul, F. and Postlewaite, A. Asymptotic efficiency in large economies with asymmetric information. *Econometrica*, 60:1273–1292, 1992.

Haas, C., Kimbrough, S. O., and van Dinther, C. Strategic learning by e-service suppliers in service value networks. *Journal of Service Research*, 16(3):259–276, 2013.

Hang, C.-W. and Singh, M. P. From quality to utility: Adaptive service selection framework. In *Proceedings of the 8th International Conference on Service-Oriented Computing (ICSOC 2010)*, pages 456–470, San Francisco, California, USA, 2010.

Harris, M. and Raviv, A. A theory of monopoly pricing schemes with demand uncertainty. *American Economic Review*, 71:347–365, 1981.

Harris, M. and Townsend, R. M. Resource allocation under asymmetric information. *Econometrica*, 49:33–64, 1981.

Harsanyi, J. C. Games with incomplete information played by Bayesian players. *Management Science*, 14:159–182, 320–334, 486–502, 1967-1968.

Holmström, B. and Myerson, R. B. Efficient and durable decision rules with incomplete information. *Econometrica*, 51(6):1799–1819, 1983.

Huget, M.-P. and Koning, J.-L. Interaction protocol engineering. In Huget, M.-P., editor, *Communication in Multiagent Systems, Agent Communication Languages and Conversation Polocies*. Springer, 2003.

Johnson, T. Matching through position auctions. Working Papers 001, Department of Economics, University of Notre Dame, 2011.

Johnson, T. Matching through position auctions. *Journal of Economic Theory*, 148: 1700–1713, 2013.

Karaenke, P. *Multiagent resource allocation in service networks*. PhD thesis, University of Hohenheim, 2014.

Korb, K. B. and Nicholson, A. E. *Bayesian Artificial Intelligence*. CRC Press Inc., 2010.

Krishna, V. *Auction Theory*. Academic Press, San Diego, 2002.

Kultti, K., Mauring, E., Vanhala, J., and Vesala, T. Adverse selection in dynamic matching markets. *Bulletin of Economic Research*, 67(2):115–133, 2015.

Kundra, V. Federal cloud computing strategy. Technical report, The White House, Washington D.C., 2011.

Kwasnica, A. M., Ledyard, J. O., Porter, D., and DeMartini, C. A new and improved design for multiobject iterative auctions. *Management Science*, 51(3):419–434, 2005.

Laffont, J. J. and Maskin, E. A differentiable approach to expected utility maximizing mechanisms. In Laffont, J. J., editor, *Aggregation and Revelation of Preferences*. North-Holland, 1979.

Lai, K., Rasmusson, L., Adar, E., Zhang, L., and Huberman, B. A. Tycoon: an implementation of a distributed, market-based resource allocation system. *Multiagent and Grid Systems*, 1(3):169–182, 2005.

Lamparter, S., Ankolekar, A., Studer, R., and Grimm, S. Preference-based selection of highly configurable web services. In *Proceedings of the 16th international conference on World Wide Web (WWW 07)*, 2007.

Lehmann, D., O'Callaghan, L., and Shoham, Y. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):577–602, 2002.

Leyton-Brown, K. and Shoham, Y. Mechanism design and auctions. In Weiss, G., editor, *Multiagent systems*. MIT Press, 2013.

Lindgrena, I. and Jansson, G. Electronic services in the public sector: A conceptual framework. *Government Information Quarterly*, 30(2):163–172, 2013.

Mas-Colell, A., Whinston, M. D., and Green, J. R. *Microeconomic Theory*. Oxford University Press, New York, 1995.

MATLAB. *Version 8.6.0 (R2015b)*. The MathWorks Inc., Natick, Massachusetts, 2015.

McLean, R. and Postlewaite, A. Informational size and incentive compatibility. *Econometrica*, 70(6):2421–2453, 2002.

McLean, R. and Postlewaite, A. Informational size and efficient auctions. *Review of Economic Studies*, 71:809–827, 2004.

Milgrom, P. and Weber, R. A theory of auctions and competitive bidding. *Econometrica*, 50:1089–1122, 1982.

Milgrom, P. *Putting Auction Theory to Work*. Cambridge University Press, 2004.

Milgrom, P. and Roberts, J. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, 58(6):1255–1277, 1990.

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

Mookherjee, D. and Reichelstein, S. Dominant strategy implementation of bayesian incentive compatible allocation rules. *Journal of Economic Theory*, 56(2):378–399, 1992.

Muthoo, A. and Mutuswami, S. Competition and efficiency in markets with quality uncertainty. Economics Discussion Papers 593, Department of Economics, University of Essex, 2005.

Muthoo, A. and Mutuswami, S. Imperfect competition and efficiency in lemons markets. *The Economic Journal*, 121(552):580–594, 2011.

Myerson, R. Optimal auction design. *Mathematics of Operation Research*, 6:58–73, 1981.

Myerson, R. B. Incentive compatibility and the bargaining problem. *Econometrica*, 47:61–73, 1979.

Myerson, R. B. and Satterthwaite, M. A. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29:256–281, 1983.

Nisan, N. and Ronen, A. Algorithmic mechanism design. *Games and Economic Behavior*, 35:166–196, 2001.

Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V., editors. *Algorithmic Game Theory*. Cambridge University Press, Cambridge, England, 2007.

OASIS. Advancing open standards for the information society, 2016. URL `https://www.oasis-open.org`. [Online; accessed 06-July-2016].

O'Sullivan, J., Edmond, D., and Hofstede, A. T. What's in a service? Towards accurate description of non-functional service properties. *Distributed and Parallel Databases*, 12:117–133, 2002.

Pal, R. and Hui, P. Economic models for cloud service markets. In Bononi, L., Datta, A., Devismes, S., and Misra, A., editors, *Distributed Computing and Networking*, volume 7129 of *Lecture Notes in Computer Science*, pages 382–396. Springer Berlin Heidelberg, 2012. ISBN 978-3-642-25958-6.

Papazoglou, P. *Web Services: Principles and Technologies*. Prentice Hall, 2008.

Parsons, S. and Wooldridge, M. Game theory and decision theory in multiagent systems. *Autonomous Agents and Multi-Agent Systems*, 5(3):243–254, 2002.

POPCORN project, 1998. URL `http://www.cs.huji.ac.il/~popcorn/index.html`. [Online; accessed 12-February-2016].

Regev, O. and Nisan, N. The popcorn market - online markets for computational resources. *Decision Support Systems*, 28(1-2):177–189, 2000.

Rust, R. and Kannan, P. *E-Service: New Directions in Theory and Practice*. ME Sharpe, 2002.

Rustichini, A., Satterthwaite, M., and Williams, S. Convergence to efficiency in a simple market with incomplete information. *Econometrica*, 62(5):1041–1063, 1994.

Samimi, P., Teimouri, Y., and Mukhtar, M. A combinatorial double auction resource allocation model in cloud computing. *Information Sciences*, 357:201–216, 2016.

Samuelson, W. Bargaining under asymmetric information. *Econometrica*, 52(4): 992–1005, 1984.

Sandholm, T. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135:1–54, 2002.

Satterthwaite, M. and Williams, S. *The Bayesian Theory of the k-Double Auction*, chapter 4, pages 99–123. Addison-Wesley, 1993.

Schnizler, B., Neumann, D., Veit, D., and Weinhardt, C. Trading grid services - a multi-attribute combinatorial approach. *European Journal of Operational Research*, 187:943–961, 2008.

Sen, A. K. *Collective Choice and Social Welfare*. Holden-Day, 1970.

Shimer, R. and Smith, L. Assortative matching and search. *Econometrica*, 68(2): 343–369, 2000.

Stösser, J., Neumann, D., and Weinhardt, C. Market-based pricing in grids: On strategic manipulation and computational cost. *European Journal of Operational Research*, 203(1):464–457, 2010.

van der Vaart, A. W. *Asymptotic Statistics*. Cambridge University Press, 2000.

Varian, H. R. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2007.

Vickrey, W. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

Waldspurger, C. A., Hogg, T., Huberman, B. A., Kephart, J. O., and Stornetta, W. S. Spawn: A distributed computational economy. *Software Engineering*, 18(2):103–117, 1992.

Wilson, R. Incentive efficiency of double auctions. *Econometrica*, 53:1101–1115, 1985.

Zacharias, E. and Williams, S. R. Ex post efficiency in the buyer's bid double auction when demand can be arbitrarily larger than supply. *Journal of Economic Theory*, 97(1):175–190, 2001.

Zaman, S. and Grosu, D. Combinatorial auction-based allocation of virtual machine instances in clouds. *Journal of Parallel and Distributed Computing*, 73:495–508, 2013.