Subproject A3:
The Market for Services: Incentives, Algorithms, Implementation

Claus-Jochen Haake[1], Burkhard Hehenkamp[1], Gleb Polevoy[2]

1 Department of Economics, Paderborn University, Paderborn, Germany
2 Heinz Nixdorf Institute and Department of Computer Science, Paderborn University, Paderborn, Germany

1 Introduction

The arguably widest definition of a **market** is the *organization of interaction across several agents*. The term *interaction* can thereby be filled in different ways, meaning that the market participants face various strategic possibilities to take action. The goal of Subproject A3 "The Market for Services: Incentives, Algorithms, Implementation" is to take a close look at interactions within and between the different groups of agents in an OTF market with the aim to elicit incentive structures, to evaluate the outcomes, or to design rules according to which the agents are supposed to act. Analyzing given interaction forms either means identifying equilibria, which enables us to predict the behavior of the market participants and thus gives us a tool to assess the consequences on individual and social welfare.

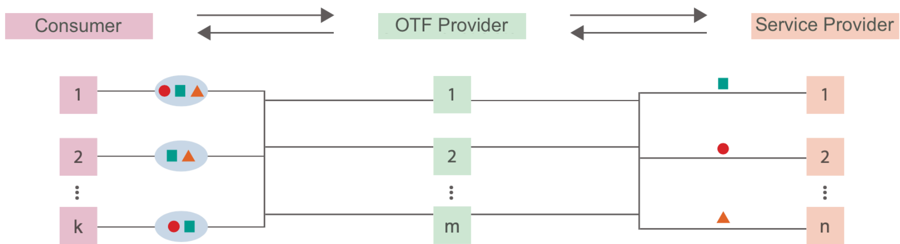

Figure 3: *Layout of an OTF market.*

The general layout of an OTF market is illustrated in Figure 3, which shows a number of specific characteristics. First, we may (ideally) view it as a two-sided market with consumers (customers, end users) one one side and service providers on the other side. OTF (service) providers act as mediators between consumers and service providers. However, unlike in traditional models for two-sided markets, OTF service providers do more than just mediate. Based on a customer's request, they actively form service compositions for which they act as a demander of services themselves. In relation to the customer, an OTF service provider takes the role of a seller to satisfy th customer's demand. As a result,

cjhaake@wiwi.upb.de (Claus-Jochen Haake), burkhard.hehenkamp@upb.de (Burkhard Hehenkamp), gpolevoy@mail.upb.de (Gleb Polevoy)

contracts are not only signed between the customer and the OTF provider, but also between the OTF service provider and possibly several service providers.

Within this market structure we investigate several different forms of interaction using methods from microeconomics as well as from non-cooperative, cooperative, and algorithmic game theory. Interaction takes place on different levels, i.e., involving smaller or larger groups. Identifying the agents' incentives and the design of institutions that lead them in the right direction is at the heart of Subproject A3.

As an example, we consider the contracting problem between a single OTF service provider and a service provider. On this micro-level, the question is how to find a fair and efficient way to settle negotiations on the terms of trade. The models we use here stem from cooperative game theory or, more specially, from bargaining theory. It turns out that there is a tradeoff between fairness and efficiency of the outcome. While the former is attractive to market participants, the latter property is, given an appropriate business model, interesting to the OTF market provider.

Taking a slightly broader perspective, an antecedent problem occurs: namely, who contracts with whom in the market. Such problems are central in matching theory, which designs and studies mechanisms that match agents from two different sides (in a market) in a preferably efficient and stable manner. Especially when capacity constraints (of OTF providers) play a role, the answer to the question of which end users should be matched to which OTF service providers could have a crucial impact on market performance. Further, matching mechanisms themselves provide incentives to act strategically and thus have to be analyzed and controlled in this direction.

From an overall view, competition naturally is a vital issue and raises a couple of questions such as how prices are formed or how the structure of offered services evolves. Concerning the latter, there may arise incentives to offer particular services only in combination with other services (bundling). From a mechanism design point of view, welfare analyses are inevitable to find directions in which the market should be regulated.

Finally, the success of an OTF market is ultimately linked to information on services that is as clean as possible. While from the outset, incomplete information on service qualities is a typical characteristic of an OTF market, the design of information systems that are based on user ratings help to reduce the lack of information, reduce incentives for misbehavior and hence increase efficiency of the market outcome. Again, on the one hand it is our task to investigate the participants' incentives to react to a well-functioning rating system that aims at distinguishing high and low quality services. On the other hand, the challenge is to elicit and process information on observed product qualities so that unwanted behavior can readily be detected.

In the next section, we highlight contributions from Subproject A3 concentrating on economic issues of the market. Abstracting from technical aspects of service composition, service execution, the organization of the infrastructure, or security issues, our task is to analyze the behavior of market participants from a mostly theoretical perspective, including the provision oof methods for efficient computation of solutions.

2 Highlights and Lessons Learned

This section is structured in three subsections highlighting our results. Subsection 2.1 focuses on specific scenarios in which (OTF) providers compete with each other and investigates the effects on welfare and market structure. In Subsection 2.2 we are concerned with a more individual level of interaction and implementation of outcomes. Further, we take a closer look at how quality standards can be maintained through reputation systems. Finally, Subsection 2.3 highlights our results on algorithms developed to calculate and hence predict market outcomes such as equilibria.

2.1 Competition

We take four different perspectives in our analysis of competition in OTF markets. First, we investigate the incentives to combine elementary products to compositions of products rather than to sell them as separate entities.[1] Second, we study experimentally to what extent implementing competition can help improve the quality of services in (OTF) markets that are characterized by fixed or regulated prices. Third, we explore whether opening a monopolistic (OTF) market for competition improves market performance by increasing the quality of products when the market entrant has the option to differentiate its product away from the already existing one. Fourth and finally, we examine how competition for innovation is affected when firms struggle for their survival in the market.

Bundling
In our first highlight, we explore a question that lies at the heart of any OTF market [EHH22]. Under which conditions is the composition of products, i.e., *bundling*, optimal to an OTF provider at all? To what extent does the incentive to sell products in bundles rather than in separate entities depend on the nature of the original elementary products such as their degree of product differentiation. Does the incentive depend on whether the elementary products represent substitutes or complements? Does it depend on their respective degrees of substitutability or complementarity?

To investigate these issues we use a rather specific asymmetric market setup, which we describe in detail further below. Additionally, we show that this specific market setup may arise endogenously in a richer setup, where service providers are enabled to choose their distribution channels optimally. Finally, we also take a social perspective and examine the welfare consequences of bundling in our framework.

Figure 4 below illustrates our market setup. We consider two retailers R_A and R_B and two monopolistic service providers M_1 and M_2, the latter of which each produce a differentiated elementary product. The retailers compete in prices. Service provider M_2 sells its product to both retailers, while M_1 only supplies retailer R_A. Retailer R_A hence receives both elementary products and considers whether or not to sell them as a bundle or as separate entities.

Accordingly, retailer R_A can be viewed as a firm deciding whether or not to adopt the role of an OTF provider. To examine retailer R_A's incentive for bundling, we determine how

[1] While market participants may compete in products or services (or both) in most of our projects, we shall henceforth write either 'products' or 'services' implicitly including the other meaning(s) as well.

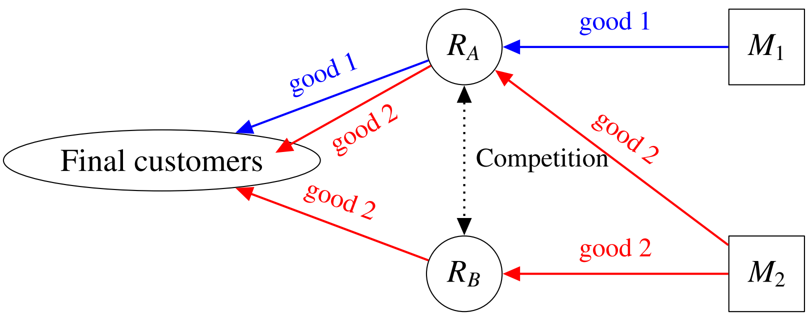

Figure 4: *The market setup.*

the degree of product differentiation of the elementary products affects the equilibrium prices, quantities and profits for each of the two selling alternatives: separate selling and bundling.

Our main result is the following: Retailer R_A will only find it optimal to bundle the elementary products when they represent close substitutes. This might sound counterintuitive at first, but we can frequently observe this form of bundling in the real world. For instance, grocery stores sell packs of peppers in bundles, either as almost perfect substitutes of identical color or as close substitutes in different colors. Similarly, clothing shops commonly offer bundles of socks or pants that only vary in patterns or colors from each other.

The intuition of our first result runs as follows. When the elementary products are differentiated, bundling alleviates competition in two ways. On the one hand, it eliminates the perfect substitutability of the elementary product 2, which is available from both retailers. On the other hand, bundling also creates a complementarity between the elementary products 1 and 2. Competition between differentiated products is the stronger the lower their degree of differentiation. Consequently, the anticompetitive effect of bundling is strongest when the elementary products constitute close substitutes, since then competition is intense. Only in this case, is the competition-reducing effect of bundling so strong that it outweighs the aggravation of the double marginalization problem that occurs along the vertical supply chain. Therefore, bundling by retailer R_A in our market setup is only profitable when the elementary products represent close substitutes. Only then might retailer R_A actually adopt the role of an OTF provider.

With regard to social welfare, the picture is blurred. Product bundling reduces the consumer surplus because of the higher downstream prices, while it increases the producer surplus, since all firms earn a higher profit, both in the downstream and the upstream market. Social welfare, however, only increases through bundling if the elementary products constitute close to perfect substitutes.[2] Thus, the emergence of OTF providers must be viewed quite critically from a social perspective.

Finally, we extend our framework to incorporate the service providers' choice of their distribution channels. As it turns out, our asymmetric market setup indeed represents an equilibrium outcome of this extended model when the elementary products represent close but not too close substitutes. Moreover, in this case, bundling reduces social welfare and should be prohibited.

[2]Recall that social welfare is defined as the sum of consumer surplus and producer surplus.

To sum up, our first highlight stresses the importance of product differentiation for bundling elementary products when the associated vertical market structure is prone to double marginalization. At least within our market framework, product bundling, and hence OTF providers, should raise serious antitrust concerns.

Introducing Competition I

Many healthcare markets can be viewed as instances of two-sided OTF markets with indirect network externalities. For example, hospitals match patients with various types of diseases with doctors from various fields of specialization. The more fields of specialization a hospital covers, the more likely it is that a patient will find the appropriate treatment. The more patients there are, the more likely it is that a doctor of a certain specialization will find patients.[3] Similarly, patients attending a general practitioner (GP) benefit from the experience the GP has gained from treating other patients. The more other patients there are, the better the experience of the GP.

This second highlight and the third one further below both examine monopolistic provider markets that are opened for competition. The aim is to evaluate whether opening the market for competition improves the market outcome in terms of quality, customer benefit or social welfare. Furthermore, we study the consequences of the providers' customer orientation (such as physician altruism) for the market outcome. In this second highlight, our focus is on the quality of the service, while the type of service is taken as given. Moreover, we deploy an experimental approach to study the effect of introducing competition. In contrast, the third highlight further below scrutinizes a theoretical model, where the new provider chooses its quality of service but may also adjust the type of service offered in order to soften competition.

Previous research has shown that, without competition, providers deviate from the customer-optimal provision under payment systems such as capitation and fee-for-service. While capitation corresponds to a fixed payment per treatment, the total payment under fee-for-service depends on the number of services executed within a given treatment. Correspondingly, a profit-maximizing provider would execute no services at all under capitation, while it would perform the maximum number of feasible services under fee-for-service. While competition is expected to mitigate these distortions, providers usually interact with each other repeatedly over time and only a fraction of customers switches providers at all. Both features might prevent the desired effect from introducing competition.

We consider two setups ([BHK17], [BHK23]). In both setups, we experimentally study the effect of introducing competition among providers when there is a trade-off between the choice of maximizing customer utility and the choice of maximizing a provider's profit. While in [BHK17] providers face homogenous customers that all face a problem of identical severity, [BHK23] scrutinizes the effects that originate from a heterogenous customer population. For both setups, we develop a theoretical model that serves as our benchmark, which we then test in a controlled laboratory experiment.

In [BHK17], our experimental conditions vary the physician payment scheme (capitation vs. fee-for-service) and the severity of the patient's problem (high vs. low). Real patients benefit from the provider decisions made in the experiment. We find that, in line with

[3] Notice that there might be also indirect network externalities, which are negative. For instance, the number of nurses per patient is decreasing in the number of patients treated.

the theoretical prediction, introducing competition can reduce underprovision (under capitation) and overprovision (under fee-for-service). The strength of the observed effects, however, depends on the severity of the problem and the payment scheme. We also find providers to collude tacitly, in particular under fee-for-service payment. Collusion appears less often than in related experiments on price competition though.

In [BHK23], providers face heterogenous customers that differ in the severity of their problem and in their mobility. Mobile customers choose their provider on the basis of the (expected) benefit from treatment, while immobile patients always visit the same provider. While we also examine the effect of introducing competition, the analysis of the second setup centers on the effects of customer heterogeneity on the market outcome. In line with the theoretical prediction, we find that introducing competition significantly increases patient benefit for mobile patients. In contrast, for immobile patients, competition worsens the outcome compared to a situation without competition. This latter observation does not match with our theoretical prediction, which would predict no difference. With repetition of the interaction both effects become more pronounced. Our results imply that introducing competition does not entail unique positive effects, but rather ambiguous effects that differ across customer groups. In particular, customer mobility is decisive for the market outcome.

Introducing Competition II

In the second highlight, we evaluate whether opening a formerly public (or private) price-regulated monopoly market for competition represents a viable option for improving quality and choice for customers. To this end, the welfare effects from opening the market are determined.

In price-regulated monopoly markets we typically observe low product quality. In principle, opening the market for competition could be a good idea if the entrant(s) offered the same or a similar product. Then, quality competition would be intensified and firms would offer higher levels of quality to attract demand. On the other hand, entrants face an incentive to avoid or at least soften competition by offering a product that differs sufficiently from the original one. As a consequence, quality might not increase to the extent expected in the first place.

To explore our research question, we consider a three-stage duopoly model of location choice and quality competition with price regulation and costly relocation. There are three active players: a budget-constrained regulator, the incumbent monopolist, and the entrant. At stage 1, the regulator sets the price. At stage 2, the entrant chooses a location, while the incumbent monopolist is already located at the center of the market and it is too costly for him to relocate. At stage 3, the two firms compete in quality for customers' demand. Observe that the location choice of the entrant at stage 2 exhibits the trade-off between moving away from the incumbent to soften competition (the so-called *competition effect*) and moving closer to steal demand (the *demand effect*).

We consider two setups, both of which are relevant in markets such as public health care, education and schooling, or postal services ([HK20], [HK23]). In [HK20], the incumbent monopolist represents a public provider, e.g., a hospital, with some degree of customer-orientation, for example, since the hospital's physicians show some degree of altruism towards their patients. In addition, the public provider faces a profit constraint that prohibits losses. The entrant, in contrast, maximizes pure profit. In [HK23], both the incumbent monopolist and the entrant are for-profit providers, each one maximizing its own profit.

Our main results are as follows: In [HK20], opening a public hospital market typically raises quality. The private provider strategically locates towards the corner of the market to avoid (too) costly quality competition. The consequences for social welfare depend on the size of the regulator's budget and on the degree of customer orientation of the public provider. If the regulator's budget is large, high quality is implemented and welfare is highest in a duopoly whenever entry is optimal. However, when the budget is small, quality levels in the duopoly correspond to the monopoly level. In particular, it turns out that for intermediate budgets it can be optimal for the regulator to not use the entire budget.

In [HK23], the entrant strategically locates towards the corner of the market, keeping the incumbent at the monopoly quality level when the regulated price is low or intermediate. In this case, quality is only raised for the entrant's customers. When the price is high, the entrant locates at the corner of the market and both providers implement a higher quality compared to the monopoly level. Moreover, the entrant always implements a higher quality than the incumbent provider. Social welfare is always higher in a duopoly if the cost of quality is low. For higher levels of quality cost, welfare is non-monotonic in the price. Therefore, the regulator will optimaly withhold part of its budget for certain budget sizes. Finally, the welfare effect from opening the market for competition depends on the price and the size of the entry cost and the decision to allow entry should be conditioned on an assessment of the entry cost.

Firm survival and innovation

Taking a dynamic perspective on competition, not only the existing products matter for the outcomes of competition in a market but also the incentives to come up with new products, new services or new technologies. This is where our last highlight departs [GHL19]. We want to explore the outcomes of competition in innovation contests where a finite number of firms potentially compete with each other for an innovation. In particular, we want to investigate the case in which the number of firms actually competing for the innovation is uncertain and in which the behavior of a firm is governed by the strive for survival of the firms.

Competition for innovation can assume different forms. For instance, there are prizes announced for certain ideas or solutions, e.g., for algorithms that manage to accomplish a certain task, or there always is the option to register an innovation as a patent. The monopoly right that comes with a patent can then subsequently be exploited by marketing the innovation as a new product or service.

Competition for innovation is characterized by three unique features, the combination of which distinguishes it from competition in product markets. First and foremost, competition for innovation is dynamic, which makes it necessary to incorporate a time dimension in the analysis. Second, the success of investments in innovation is highly uncertain and depends on the investments made and the outcomes realized by a firm's competitors. Moreover, not always the firm with the highest investment will succeed in winning the patent. Third and finally, the investments of *all* competitors are sunk, also those made by the unsuccessful firms. Accordingly, a dynamic contest model with imperfect discrimination represents the appropriate model to study our questions at hand.

To model uncertainty about the number of competitors, we investigate two setups. In the first, we consider stochastic participation, that is, a single firm does not know the number of competitors that also participate in the same contest. From the perspective of the firm, it appears *as if* other firms participate in the contest with a certain probability. In addition, the number of potential competitors in the contest is exogenously given. In the second setup, we depart from the latter assumption to determine the number of potential competitors endogenously. To this end, we introduce a fixed cost of entry.

Markets with a high innovation intensity tend to be very dynamic and subject to high fluctuations. The entry and exit of firms are the rule rather than the exception. Correspondingly, the behavior of a firm is better described as governed by the firm's striving for survival than by profit maximization. To account for this, we deploy the so-called economic evolutionary approach.

Under the economic evolutionary approach, the biological evolutionary forces of selection, mutation, and heredity correspond to economic evolutionary forces such as imitation, innovation, and bankruptcy to name but a few. (Economic) evolutionary equilibrium then serves as the short-cut to the evolutionary outcome of a dynamic evolutionary process of imitation and innovation. Finally, we consider finite (firm) populations as the economically relevant case and solve our model for the corresponding equilibrium of a finite-population evolutionarily stable strategy.

Apart from solving for the economic evolutionary equilibrium, our focus is on the *issue of (over-)dissipation.* This issue is closely related to Posner's (1975) famous full dissipation hypothesis, according to which competition for a monopoly (in our case: a patent monopoly) would eat up the entire monopoly rent that the firms compete for. In our paper, we re-evaluate Posner's hypothesis. In our setup of a finite population, the strive for survival leads to more aggressive investment behavior, so the issue might be particularly pronounced.

Our main findings are as follows. Firstly, when the probabilities of participation are exogenously given, competitors choose higher levels of investment in the economic evolutionary equilibrium than in the Nash equilibrium. Moreover, there is ex-ante overdissipation in the economic evolutionary equilibrium for sufficiently large probabilities of participation if, and only if, the impact function is convex.[4] These results generalize earlier findings in [HLP04] from contests with a given, i.e. deterministic number of firms to contests with stochastic participation, where the number of actual competitors is *a priori* uncertain.

Secondly, with costly endogenous entry, firms enter the contest with a higher probability and choose higher levels of investment in the economic evolutionary equilibrium than in the Nash equilibrium. Importantly, under endogenous entry, overdissipation can occur for all types of contest technologies, in particular those with concave impact functions.

Our findings point to potentially high welfare losses stemming from innovation contests when they are open to anyone.

[4]An *impact function* can be understood as a lottery production function, which governs the transformation of investments into probabilities of winning. A *convex* / *linear* / *concave* impact function then exhibits increasing / constant / decreasing returns to investments.

2.2 Allocation and Incentives

In the OTF market interaction takes place in various environments. Especially when there are few participants, competition may not be a plausible form of modeling market organization. Therefore, we analyze alternative models of interaction such as bargaining, matching or, more generally, mechanism design. Because there is naturally no complete information on service qualities, we need to keep track of the agents' incentives to misbehave and exploit a superior information position at the cost of efficiency. In this subsection we review some of our results in the above fields.

Bargaining

The term *bargaining* generally refers to a situation, in which two or more persons can sign an agreement on the distribution of joint gains. A *bargaining solution* is supposed to propose such an agreement for any possible bargaining problem (from a specific set of problems), i.e., not focusing on a particular problem. The main challenge is to define or design bargaining solutions that capture the context of interaction and in which agreements are to be considered as *fair*. In this section, we briefly review two works that argue for the selection of particular bargaining solutions in bargaining problems occurring on an OTF market.

Some particular interactions in the OTF market involve two participants only. For instance, consider the situation, in which an OTF provider and a service provider negotiate over the terms of trade, including the price. Since we consider highly specialized services, there is typically no market, in which a price can be settled by an equilibrium mechanism, which means it is subject to bilateral negotiations. To aggravate the problem, both parties may have private information on either production costs or expected revenues from sales to the customer. In a simplified version, the OTF provider and the OTF service provider negotiate over a service level (e.g., quantity, quality degree, etc.) and a total payment. The presence of incomplete information requires the two parties bargain over contracts that are type dependent, i.e., those that depend on the realization of costs and revenues that have to be reported by the parties. Because the final agreement (payment, quantity, etc.) relies on the report of unobservable private information, it is naturally open to cheating. As a result, a "good" contract should satisfy a number of properties such as (i) ex post Pareto efficiency (EPE), (ii) individual rationality (IR), and (iii) incentive compatibility (IC). While EPE ensures that there is no room for renegotiations after the contract has been agreed upon, IR provides incentives to enter negotiations at all. Finally, IC requires that truthful reporting be a Nash equilibrium in the reporting game.

A closer inspection of the structure of the negotiation problem reveals that it exhibits the features of an intra-firm transfer pricing problem. A transfer pricing problem as it is studied in the literature involves two divisions of the same company and is displayed in Figure 5. The aim is to settle an agreement on how much and at what price an intermediary good is internally sold from the producing division to the buying division. While the producing division has better information on production costs, the sales division knows the external market on which the product is finally sold to customers. Besides the goals of the divisions that act as profit centers, the central management prefers a company-optimal outcome. Returning to our negotiation in the OTF market, the additional quantity parameter may be interpreted as such if a hardware service is under concern. Alternative interpretations may include the duration of service provision. The role of the company could be played by an OTF market maker who is interested in an efficient market outcome.

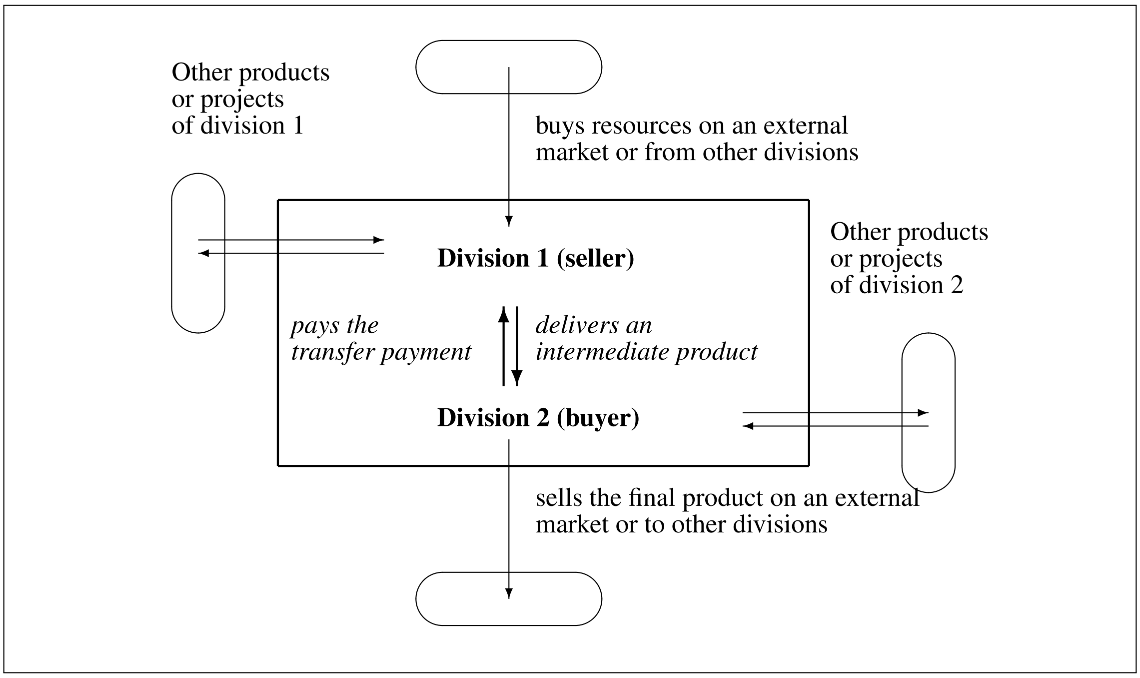

Figure 5: *Intra-firm Transfer Pricing Problem. The two divisions bargain over a payment from division 2 to division 1 in return for a quantity of the product. (Source: [HR18]).*

In [HR18] we propose a fair solution to the transfer pricing problem that rests on a solid foundation in bargaining theory and which is new to the transfer pricing literature. We mainly use the transfer pricing scenario as we compare the solution to other well-known solutions from this field. To be more precise, for a transfer pricing game under incomplete information we determine the generalized Nash bargaining solution. Requiring agreements to be incentive compatible and/or efficient, we further highlight the relation between these two desirable properties. For a necessary intermediate result for the applicability of the generalized Nash bargaining solution, we show that the transfer pricing game is regular, meaning that it is possible to guarantee each division a strictly positive expected profit, regardless of their specific private information. From a managerial perspective, the appealing feature of the generalized Nash bargaining solution is that it provides each division with a strictly positive expected profit. Further, we derive necessary conditions for a mechanism that implements the generalized Nash bargaining solution (Propositions 4, 5, and 6) and shed light on the trade-off between efficiency and fairness (Proposition 7). As illustrated in examples, the Nash solution tends to keep differences in divisional profits smaller in comparison to other solutions. Two examples illustrate differences between the generalized Nash bargaining solution and well-established alternatives from [Wag94].

In sum, we find that if parties are interested in a fair outcome, our analysis provides good arguments to use the generalized Nash bargaining solution. For the bargaining problem on an OTF market, we may think of bargaining problems as being automatically resolved in a way that takes fairness and efficiency into account. While the former increases the attractiveness of an OTF platform, the latter increases the incentives of a market maker to provide the platform itself.

The nature of a (cooperative) bargaining problem is to distribute common gains in a fair way. Apart from bargaining particular bargaining situations in which parties bargain once, we also investigated problems of repeated interaction from a structural point of view. [Hoo20] analyzes a model in which the bargaining problem is shaped by the possible payoffs from strategies in a differential game, i.e., a non-cooperative game that is played in continuous time. Thus, in this scenario the participants may share common gains over time. In terms of fairness, this requires that solutions be individually rational over time and consistent with time preferences, which can be thought of as discount factors for future payoffs. The main result is that for a class of underlying differential games, both properties are already fulfilled when the bargaining solution satisfies an overall individual rationality. One advantage of the latter property is that it is intuitive, as it guarantees participation in the interaction and in accordance with the main result, it triggers consistency over time. From a theoretical point of view, this has an impact on which bargaining solution should or should not be selected by a market designer.

Matching

The problem of finding a good market allocation is directly connected to the question, who serves whom in the market. There is a still growing literature on two-sided matching markets that discusses algorithms for matching agents from one group to agents of another group. The assessment whether an algorithm is “good” or “bad” is ultimately linked to the properties of the final outcome. Besides the efficiency, stability of the matching plays the most important role. It guarantees that no agent or group of agents would want to alter the matching and have the possibilities to do so. Examples for classical matching markets are the marriage market and college admissions [GS62], school choice [AS03] and the housing market [SS74].

In an OTF market, specific allocation problems can be viewed as a matching market – end users have to be matched to OTF service providers, while service providers are matched to OTF service providers. The ingredients of a matching market are the participants’ preferences over participants on the other side. For example, differences in the characteristics of an OTF provider or the heterogeneity of traded composed services form a consumer’s preferences over providers, while the users’ different demands or their willingness to pay shape an OTF service provider’s preferences over users. Regarding an OTF market as being organized on a (central) platform operated by an OTF market provider, the matching problems can be described as a many-to-one matching market, which in the literature is commonly termed the *college admission problem (CAP)* or school choice problem.

We address three main problems connected to desirable matchings in the OTF market: (1) How can matching algorithms be adjusted to cope with users’ heterogeneous demands? (2) Since the market is large, a participant may not know all options on the other side, but still has to form preferences. Therefore, how can incomplete information on the participants be dealt with? (3) Finally, how is the functioning of matching mechanisms affected, when the formation of preferences follows an (automatic) pattern?

Question 1 addresses a problem that is widely ignored in classical matching models. In a CAP, students are matched to exactly one seat at a college, so that all have the same weight or need of capacity. The total number of available seats, or total capacity, belongs to a college's characteristics ([GS62; AS03]). For this scenario, the Boston Mechanism (BM), the Deferred Acceptance Algorithm (DA), and the Top Trading Cycle Algorithm (TTC) are the most used mechanisms in practice. But if the homogeneous demand assumption that each student requires exactly one seat is dropped, only little is known. In weighted matching markets or matching markets with sizes [BM14], stability can no longer be assured [MM10].

In [HS20b] we start from the fact that stability is no longer assured and investigate how we can find a stable outcome, if possible, and how to enable it otherwise. For this, we introduce a new algorithm, the deferred acceptance algorithm with gaps, which either results in a stable matching, if one exists, or leads to a cycle. If all students have the same weight, meaning that they all need the same amount of capacity, the algorithm operates exactly as the deferred acceptance algorithm [GS62]. However, if the algorithm leads to a cycle, because there is no stable matching, we can arrive at stability by increasing or decreasing the colleges' capacities.

As stability is no longer guaranteed unless we modify capacities, it is also possible to have a look at Pareto efficient outcomes of weighted matching problems or, more precisely, weighted school choice problems. [Str20] proposes a variant of the TTC algorithm: namely, the weighted TTC (WTTC), which is strategy-proof and yields a Pareto efficient outcome. But although the main results carry over compared to the TTC, the usage of the WTTC introduces a trade-off between weights and preferences or priorities. Thus, the introduction of weights comes with some costs as it is more complex to guarantee each student a seat at a college.

Addressing the second question, one source of incomplete information over the other participants' preferences comes from the bare size of the OTF market. Incomplete information was introduced by [Rot89]. Given a strategy-proof matching algorithm, he shows that truth-telling is still a dominant strategy if agents only have limited information about the other agents' preferences. While there is already some experimental research on the functioning of algorithms that are not strategy-proof in a setting of incomplete information [CS06; CLS16], no empirical evidence exists on the behavior of students in a school choice problem with incomplete information when a variant of the Boston school choice mechanism is used. The BM does not yield a stable or Pareto efficient outcome and, most importantly, is not strategy-proof [AS03].

[HS20a] fills this lack and further introduces heterogeneity in the market by taking into account different weights of students. Our research uses two different data sources, the data derived from the clearinghouse that implements the matching algorithm itself and data from a voluntary survey among the students who participated in the clearinghouse. We find that over 74% of students misrepresent at least one of their ranks in the preference list, which is not surprising given that the algorithm is not strategy-proof. But although students are trying to exploit incentives, they do not necessarily succeed in improving their outcomes through manipulation. This is mainly due to the fact that students have incomplete information on the other students' preferences. Additionally, some students do better in misrepresenting than others. We call these students sophisticated, whereas another group of students is not able to act in a consistent manner and is thus naive. This notion of sophistication is based on the theoretical definition by [PS08]. We see that sophisticated students actually reach significantly better outcomes than naive students.

Another source of incomplete information may arise on the other market side, which brings us to the third question. In school choice it is assumed that schools are not able to rank all the students individually but with the help of some objective criteria [AS03]. But if these criteria are based on the students' preferences this might actually undermine the functioning of any matching mechanism. More precisely, [HS18] analyzes what happens if the schools' priorities are formed in a reciprocal way, i.e., based on the students' preferences in a "first-preference-first" manner. We show that in this case the deferred acceptance algorithm, the TTC, and the Boston school choice mechanism all yield the same outcome and are thus manipulable. This means that even in otherwise strategy-proof mechanisms, it is no longer a dominant strategy to state the true preferences.

To sum up, our works provide insights on how matching processes shall be set organized in an OTF market. There, heterogeneity of agents and incomplete information on the matching problems. The results show that special care has to be taken on the design strategy proof algorithms and on the formation of preferences to exclude unwanted behavior or the emergence of misdirected incentives.

Quality assurance: Customer evaluations

Addressing the question of how a good or poor service quality can be identified, we review two papers on how customer feedback can be of assistance. One of the characteristics of almost any market and in particular of OTF markets is that information on service qualities is asymmetrically distributed. One reason for this is that services are typically experience goods. Such goods do not reveal their true qualities to the consumer prior to purchase and consumption. In that spirit, whether a particular service from a service provider or the service composition sold by the OTF service provider actually delivers the desired result to the end user can only be verified after the transaction between user and OTF service provider has taken place.

A particular consequence of the observability of service quality after purchase is that it opens the door for strategic interaction on the provider's side. Because lower service quality typically comes at lower costs, an optimal decision would be not to produce high quality services. To make the problem even more demanding, service compositions may fail to work well if there is only one single "bad" service used. A non-perfect but arguably useful instrument to inform about experienced service quality is to use consumer evaluation systems such as those introduced by online retailers. Resting on an intrinsic motivation of customers to rate products, such systems may give a valuable tool for deciding which service to request or to integrate into a composition.

Taking a theoretical perspective, we addressed two questions: First, how do providers react to rating systems in the sense that they may exploit a good reputation? Second, given information on ratings for composed services, can we use it to derive a rating of the component services?

In [MFHR18] we modify the model from [Del05] and model the situation of a service provider who strategically decides to repeatedly sell its service to customers in either high or low quality. The delivered service is rated as good or bad by the customer, who is, however, not fully able to identify the true quality. We therefore model a customer's feedback as a random variable, whose distribution depends on the delivered quality. Phrased differently, from the provider's perspective, there is a higher chance to receive a good rating when the quality is good, compared to the case in which the quality is bad. Still, even with bad quality, the provider may receive a positive rating. The collection of the three most recent ratings is taken as a proxy for the provider's reputation. The service price is modeled to be directly dependent on the reputation, more precisely, on the number of positive ratings among the three most recent ratings. When making a strategic decision, the service provider has to compare the benefits from maintaining (or building) a good reputation with a higher sales price and higher revenues with the temptation of milking a good reputation by delivering bad quality and running the risk that prices will fall in response to a decline in reputation.

We address this trade-off by analyzing the theoretical model as well and testing the result in an experiment. Theoretically, we analyze the corresponding Markov Decision Problem and demonstrate that keeping the service quality constant is an optimal strategy for the service provider. In essence, whether supplying high or low quality is optimal depends on the difference of probabilities for receiving a positive rating. The larger the difference, the more accurate the consumer's rating is. This is intuitive because increasing the probability for a good rating when the quality is actually low increases the incentives to produce low quality.

Milking behavior means that the provider delivers bad quality whenever enough positive ratings appear in the rating history in order to keep the price high. When too many negative ratings appear, a good reputation is built up by delivering good quality which comes with a higher probability for good ratings. In the theoretical model the optimal strategy is either to constantly sell good or bad quality, so that milking one's own reputation does not take place. In the experiment, however, milking behavior can be observed, meaning that subjects with a high reputation tend to produce low quality for a couple of rounds, instead of maintaining the good reputation (which would have been the optimal solution). The striking lesson that we learn from this work impacts the design of reputation systems: The better the accuracy of the system, the higher the incentives to serve the market with good quality.

The second question focuses on how much we can say about the quality of single services. In [FHSS18] we discuss how to disentangle the ratings for service compositions from m services by n consumers as follows. Starting with the collection of user ratings over compositions, there has to be an aggregation step A and a disaggregation step D to arrive a rating over services, which leaves two options: a) either we first aggregate the ratings across users, which gives us an overall rating of compositions and we can then elicit (or disaggregate) information on single service ratings, or b) we first disaggregate individual ratings to individual ratings over single services, which should then be aggregated to an overall rating of single services. Figure 6 illustrates the aggregation/disaggregation problem. The starting point (see upper left corner) is a $n \times 2^m - 1$ matrix, in which each row corresponds to a user's evaluation of the $2^m - 1$ possible service compositions. The final result is supposed to be a $1 \times m$ matrix (lower right corner) that contains ratings for the m single services.

The task is to design informative aggregation operators A_1 and A_2 as well as disaggregation operators D_1 and D_2. While from the outset it is not evident which operators fit best, one essential property we impose is that the two routes sketched above yield the same result, making the diagram in Figure 6 commutative. Moreover, anonymity requirements guarantee that no single service and no two users are treated differently. Further, it should

$$\begin{tikzcd}
\boldsymbol{\mathcal{M}\,(n \times 2^m - 1)} \arrow[r, "D_1"] \arrow[d, "A_1"] & \mathcal{M}\,(n \times m) \arrow[d, "A_2"'] \\
\mathcal{M}\,(1 \times 2^m - 1) \arrow[r, "D_2"] & \boldsymbol{\mathcal{M}\,(1 \times m)}
\end{tikzcd}$$

Figure 6: *Aggregation and disaggregation (from [FHSS18]).*

not be possible for a single user to significantly manipulate the rating of a single service by changing its own valuations.

For the disaggregation step, we reinterpret a rating over compositions as a cooperative game with transferable utility (TU game) and use the Shapley value as a solution concept to attribute a rating to each single service/player. For an aggregation device (across users), we use the averaging operator. This combination turns out to be commutative in the sense above. It is anonymous and no single user has a strong influence on the final valuations. Other (intuitive) methods such as taking minimal or maximal composition values particularly fail to satisfy this non-manipulability property.

For the OTF market, this means that a smartly designed system that processes end user valuations over composed services can help to identify those (component) services that fail to work well in compositions. This information in turn can be used by OTF service providers when composing services to satisfy a user request. Therefore, the demand for and pricing of services are influenced by the analysis of customer feedback.

Mechanism Design

Designing a market means designing the rules for interaction according to which the market participants react to and choose their strategies. The combination of strategies determine allocations, payoffs, or welfare, hence the market outcome. The design of rules such that strategic behavior finally leads to a desired (market) outcome is at the heart of implementation theory or mechanism design. At a fundamental level, the question arises which outcomes can be implemented through strategic interaction at all and what is an appropriate equilibrium concept. In [HT21] we analyze this problem in a very general model. The notion of a mechanism (describing the rules of interaction) is expanded to one of a *socio-legal system*, which allows to cope with two types of obstacles that have been widely ignored in the mechanism design literature.

First, unlike in traditional mechanism design, a player's set of feasible strategies may depend on the other players' choices of strategy. As a consequence, specific strategy profiles might not be feasible. For a simple illustrative example, consider a number of OTF service providers who choose how much capacity of a hardware resource they want to use. Because the total capacity is limited, each strategy profile of the other providers sets an upper bound for the choice of a particular provider. Phrased differently, it might occur that the total capacity chosen by all providers exceeds the maximally possible capacity of the resource. However, one of the providers could be blamed for choosing the "wrong" strategy, because feasibility is a property of the chosen profile of strategies. Second, the mechanism designer (e.g., the OTF market maker) might want to avoid "illegal" behavior in the sense that particular outcomes should not occur. As a simple example one can think of OTF service providers who can choose to either serve or not serve a particular customer. The designer can declare that all profiles in which at least one provider serves the customer are "legal" ones, because he is interested in an outcome guaranteeing that the customer is served.

The extension from mechanisms to socio-legal systems requires an adaption of the equilibrium concept. We define the notion of a *Debreu-Hurwicz equilibrium* that combines the Nash equilibrium concept with features from Hrwicz' work on legality and Debreu's social equilibrium concept (see [Hur94], [Deb52], [KY18]). The choice or design of the equilibrium concept is ultimately linked to the question which social choice rules, i.e., which desirable outcomes, are implementable.

We address this question by investigating implementability of the cooperative Nash bargaining solution, which marks the desired outcome for a population of players. We find that by gradually expanding the equilibrium notion from Nash to Debreu-Hurwicz, undesired equilibria (i.e., those that do not trigger the Nash bargaining solution as an outcome) can be removed so that we have ultimately been able to show a new uniqueness result. Although our paper does not directly construct or analyze a specific interaction in the OTF market, it aims at opening a new route in implementation theory allowing to explicitly cope with unwanted behavior by the players. This is important for the functionality of interaction.

2.3 Algorithmic Game Theory

We now review the primary works in subproject A3 that employ the algorithmic game theory approach. Game theory complements the mechanism design approach, where one designs interaction, by rather analyzing, e.g., existence and properties of equilibria in specific interactions. Algorithmic game theory studies both interactions of algorithms, such as the competition of trading or negotiation algorithms, and also algorithms executed on models of interactions, such as computing Nash equilibria or bargaining outcomes.

We start off by presenting general models of strategic sharing of resources under the umbrella of budget games, which model a market of products, and of various congestion games, which model sharing situations. We then continue with the less general but more network-specific progressive filling games, where the choice of routes determines the allocated bandwidth in a natural utility max-min fairness manner. Finally, we present a practically relevant online algorithm generalizing bin packing.

Sharing of resources
Our works contribute to the literature on sharing resources. [DRS14] study a market situation via introducing *budget games*, in which players choose tasks (products), that in turn have demands for resources. Consequently, choices have an influence on the sharing of necessary resources between chosen tasks. The budgets of resources are either shared proportionally between the tasks or dependent on the decision order. The authors studied the optimal solution, as well as the existence, complexity and efficiency of equilibria.

This model, for instance, describes resource sharing that occurs in cloud computing where the clients compete on the products of the cloud. In the strategic variant of the game, in which market entrance is simultaneous, the utility of a resource is shared proportionally. In contrast, in an ordered budget game that models the market entrance order the resources are allocated in the entrance order, and a deviator moves to the last position. Since the strategic budget game is a basic utility game, its price of anarchy is at most 2, as proven in [Vet02]. [DRS14] prove that this bound also holds for ordered budget games. First, [DRS14] prove that finding the optimal allocation is NP-hard and can be approximated within 1˘1/e, provided the players' strategies form a matroid. Concerning the Nash equilibrium, it may not exist in the strategic budget game and deciding whether it does or not is NP-hard. In the ordered budget games, even strong Nash equilibria exist and are polynomially computable. The strong price of stability is 1, while the strong price of anarchy is 2. The authors demonstrate that improvement moves converge to a Nash equilibrium, but it may take exponentially many steps.

A related model, which also studies sharing resources, albeit differently, is the one of *congestion games and their variations*, which capture many important interactions: in particular, network interactions where bandwidth, CPU or another resource is used by several parties. In the face of the need imposed by their ubiquity, computing the equilibria of congestion games is appallingly PLS-complete ([FPT04; AS08; ARV08]). Moreover, weighted congestion games may possess no potential function or even no pure Nash equilibria at all ([GMV05; FKS05]), and it is NP-hard to decide whether Nash equilibria exist ([DS08]). The only classes where a potential always exists are classes with linear or exponential cost functions ([FKS05; HK12; PS07]). But even for linear costs, computing and equilibrium is PLS-complete ([ARV08]). Since mixed equilibria are generally harder to interpret, the lack of pure ones indeed poses a problem. In order to ameliorate the existence and computation problem, [CFGS15] studies existence and structure of approximate Nash equilibria in weighted congestion games. They also proposed several algorithms to find approximate Nash equilibria. An earlier algorithm by the same authors in [CFGS11] computes a constant-approximate Nash equilibrium in *unweighted* congestion games with cost functions all being constant-degree polynomials. Another known result is that for symmetric unweighted congestion games, any $1 + \epsilon$-improvement dynamics converges to a $1 + \epsilon$-approximate Nash equilibrium in a polynomial number of steps ([CS11]). Moreover, [AAE+08] shows rapid convergence to socially efficient states, but those states need not be approximate equilibria.

[CFGS15] approximates a given weighted potential game with a special potential game termed Ψ-*game*. They approximate a weighted congestion game with cost functions of degrees at most $d \geq 2$ with a Ψ-game of degree d, and prove that a ρ-approximate equilibrium of such a Ψ-game of degree d constitutes a $d!\rho$-approximate equilibrium of the original weighted congestion game. Since the Ψ-games have potential and thus a Nash equilibrium, this implies that the original game possesses a $d!$-approximate equilibrium. They also provide polynomial approximation algorithm for constant d and bound the length of a best-response sequence from any initial state to a $d^{O(d^2)}$-approximate pure Nash equilibrium.

Following up on [CFGS15], [HKS14] sets out to improve the approximation factors of approximate pure Nash equilibria. Since the existence of an α-approximate potential function implies the existence of an α-approximate Nash equilibrium and the convergence to such an equilibrium of steps improving by the factor of at least α, they concentrate on α-approximate potential functions with smallest possible αs. For several cost functions, such as the polynomial ones or the concave ones, they prove the existence of α-approximate potential functions with smaller values of α than was previously known. Concretely, they provide the upper bounds of $3/2$ for concave cost functions and bounds of $4/3, 1.785$ and 2.326 for polynomials of degrees $2, 3$ and 4, respectively. In general, for polynomials of degree l, their bound is $l + 1$. For two players, their results are provably tight.

Progressive filling games

Congestion games constitute very important general models, but they assume a player's bandwidth is the sum of what she gets allocated on each edge, rather than the maximum thereof. This is ameliorated by the bottleneck congestion games ([CDR06]) where the bandwidth of a player is the maximum allocated bandwidth. However, the computation of an equilibrium there is NP-hard. Moreover, the main modeling disadvantage of bottleneck congestion games is the lack of flexibility in bandwidth allocation, contrary to the flexible *Max-Min Fairness (MMF)* from [BG21], which we present next. Therefore, [HHSS14] defines and analyzes *Progressive Filling Games (PFG)*, which model players choosing routes and receiving fair bandwidth according to the MMF algorithm. That work generalizes [YXF+10; YXF+13] to strong NE and a broader class of water-filling algorithms. They also provide a picture of the complexity of computing SNE and present the prices of anarchy and stability, as we now describe. MMF is a known fairness standard, where nobody's allocation can be increased without hurting a worse-off party. Some known generalizations include weighted MMF and utility MMF. This paper implements utility MMF by a polynomial water-filling algorithm. They define routing games with progressing filling, where each player picks a set of resources, aiming to optimize her allocated bandwidth. They assume that the flow control instantly converges to the corresponding generalization of MMF after each route update, an assumption justified, for example, by [WLLD05].

[WLLD05] studies the existence, the computation and the efficiency of the pure and of strong NE in these games. They first prove the existence of strong NE for any generalization of the water-filling algorithm. As long as certain conditions on the rate functions hold, conditions that cannot be dropped. The authors also suggest an algorithm to compute a strong NE, employing a packing oracle. They then present hardness results for computing strong NE. Next, the authors provide tight bounds on the prices of anarchy and stability, assuming the utilitarian social welfare, providing bounds that hold even if an arbitrary capacity-respecting allocation is allowed, not necessarily an MMF one. In general, the prices of anarchy and stability are n, and this is tight for both pure and strong equilibria. For routing a single commodity using MMF, the price of stability is $2 - (1/n)$ for both normal and strong NE, and the price of anarchy is n for NE and 4 for strong NE, all the bounds besides the latter being tight. If the allocation rules are fixed, the $2 - (1/n)$ bound cannot be overcome. However, if we can adjust the weights in the weighted MMF water-filling algorithm, then we can make the game have an optimum SNE. This is NP-hard to compute, but can be approximated.

Bin packing

Since execution of composed services is an integral part of the OTF market, an important algorithmic topic addressed in subproject A3 is cloud-server storage, balancing the server limitations and minimizing the cloud costs, including reducing the wear and tear of the server, energy costs, and the communication and the execution time. This is modeled in [FFG+18], who design an *online algorithm for dynamic bin packing* that balances competitiveness ratio with minimizing the number of repacks, called recourse. Offline bin packing is approximable with additive $O(\log OPT)$, where OPT is the optimum value, but online bin packing has a 1.540 multiplicative gap, even when the optimum value approaches infinity. This work now considers fully dynamic bin packing with bounded recourse, namely allowing arrival and departures of items and their repacking. They define worst-case and amortized recourse, measuring the movement costs at each time or in total, respectively.

This work characterizes the recourse to asymptotic competitive ratio trade-off in the following cases. For unit movement costs, they provide tight upper and lower bounds. The asymptotic competitive ratio here is better than that for online bin packing without repacking! That technique uses LPs and Myopic packing from [Ivk95]. For general movement costs, [Sei02] suggests a super harmonic algorithm with constant recourse, implying a competitive ratio of 1.589, which is close to the best known bin packing result of 1.578 [BBD+18]. Moreover, the authors conjecture that fully dynamic algorithms can be reduced to online ones, which would imply equal asymptotic competitive ratios, while maintaining a constant recourse. Finally, if the costs are just equal to the sizes, the authors provide a tight bound.

3 Conclusion and Outlook

The lessons we learned from our analyses of the OTF market and the interactions within raise new questions that are not exclusively interesting for the functioning of OTF markets. In what follows, we briefly comment on these more general implications of our findings for economic modeling and future works.

Bundling

The theory of bundling and tying is highly developed. Our paper is unique in that it specifically takes and analyzes an asymmetric exogenous market and distribution structure as a starting point. What is more, we show that under certain conditions this asymmetric distribution structure emerges as an equilibrium outcome in a richer model where the participating service providers also decide on their distribution channels.

Our question, under which conditions bundling occurs, does not only represent one of the core questions in the analysis of OTF markets, but also has important bearings for the analysis of bundling and tying in general markets. In particular, our findings point to negative welfare effects of product bundling and raise serious antitrust concerns. An empirical analysis to quantify the welfare effects would be desirable.

Opening markets for competition

Politicians and economists alike quite generally believe in the benefits of competition. Accordingly, they often advocate the opening of markets for competitors. This issue has been thoroughly analyzed for 'normal' markets, but not so much for public provider monopolies that are opened for competition to private providers. Here, the contribution of our paper sets in.

Our experimental analysis points both to beneficial effects that are caused by introducing competition, but also to negative effects that result from collusion and from the immobility (or non-responsiveness) of customers. Importantly, introducing competition amplifies treatment inequalities across customer groups.

Our theoretical analysis shows that introducing competition in a price-regulated market does not necessarily bring about higher quality. In particular, if the regulator's budget is small, opening a formerly monopolistic market for competition will entail that an entrant differentiates its product away from the existing product in order to soften competition. As a consequence quality may not be raised at all. For larger budgets, quality will eventually be raised by both the incumbent public provider and the private entrant. However, these budgets also need to be financed, the cost of which is substantial.

Our findings raise a number of interesting questions. As to our experimental analysis, it would be instrumental to identify ways that would allow to avoid the unequal treatment of customers, in particular the low quality treatment of disadvantaged ones. Our theoretical model could be extended to allow for multiple dimensions of quality (contractible vs. non-contractible) or to include horizontal dimensions other than location.

Firm survival and innovation

Our economic evolutionary approaches to the analysis of behavior in innovation contest represent the first theoretical approaches that are capable of explaining overdissipation, a phenomenon that has also been observed empirically in experiments. It would be interesting to engage in an empirical analysis that covers real markets for innovation. Here, a structural econometric approach would be appropriate.

Bargaining

The interaction among few players requires an alternative to models of competition. Instead, bargaining models are the more appropriate choice. While fairness and efficiency are compatible in the case in which players are completely informed about each other's preferences, we have identified a trade-off between these properties in the case of incomplete information. Still, in the literature on bargaining theory there are few works that actually work out this trade-off by rigorously defining and analyzing solution concepts that are adapted to the information scenario. Especially asymmetric solutions deserve a more detailed investigation, because not only in the OTF context do the players bring different skills or power to the bargaining table.

To better understand differences between the various solution concepts that are discussed in the literature, it would be worthwhile to find a common ground in the sense of a unified approach. This could be done either on a descriptive basis, so that solutions appear as special cases of a more "universal" solution, or on a normative ground, meaning that characterizing axioms are comparable. Either way, results may help to judge which bargaining solution is most appropriate for a given problem.

Matching

Matching theory is a vital and growing research field that has branched out into many diverse applications and questions. From our works on matching mechanisms we have learned that large markets, in which agents are naturally not able to know all other agents and their preferences, require a careful design of matching mechanisms that provide the right incentives.

When encountering the largeness of a market by (partly) generating preferences, particular patterns may occur that cause unwanted incentives for strategic behavior. Thus, we see a tradeoff between facilitation and prevailing incentives, which has not yet been well explored in the literature.

The treatment of primary data on preferences marks another challenge for future research. Since matching mechanisms are typically conducted by a central clearinghouse, all relevant information has to be collected at one place. A distributed version of a mechanism that collects information at different places contributes to the protection of private data and thus enhances the acceptability of the procedure. Here, cryptographic methods may help to reach decentralized versions with a minimum exchange of information.

Reputation systems

We have seen that implementing a rating system for services that is as accurate as possible reduces the incentives for providers to milk one's own reputation and deliver poor quality services at a high price. Further, a smart processing of user evaluations can identify poor quality. Both results are promising, not only in the light of OTF markets. The study of aggregation and disaggregation of user information can still be refined in a positive as well as a normative direction. From a positive viewpoint, for example, the investigation would look into how an inherent asymmetry among users (experts vs. ordinary end users) can be reflected in the aggregation process. From a normative viewpoint, a characterization of data processing through axioms is still missing with which ideally different mechanisms can be assessed by their defining properties.

Apart from theoretical work, it will be interesting to investigate how real agents react to different designs of a reputation system. As a result, experimental evidence gives (further) advice which elements of a system are important. This includes anonymity of raters, the rating scale, or (possibly aggregated) information that is displayed.

Mechanism design

Here, our research is meant to be of a more general nature. The aim was to describe a framework in which unwanted strategic behavior can be captured. In the language of OTF markets, one could say that the market provider (the mechanism designer) may use mechanisms that, e.g., rule out the use of an illegal strategy. At a conceptual level, this means that the responsibility for adhering to the rules is shifted from the players to the designer. However, this immediately renews an old question about the responsibility of the designer or, as cited in [Hur94], "but who will guard the guardians." Fitting the idea that OTF markets are self-organizing systems, one way could be to equip a mechanism with a control device that allows the players themselves to regulate the designer's choice of mechanism.

Bibliography

[AAE+08] AWERBUCH, B.; AZAR, Y.; EPSTEIN, A.; MIRROKNI, V. S.; SKOPALIK, A.: Fast Convergence to Nearly Optimal Solutions in Potential Games. In: *Proceedings of the 9th ACM Conference on Electronic Commerce*. EC '08. Chicago, Il, USA: Association for Computing Machinery, 2008, pp. 264–273.

[ARV08] Ackermann, H.; Röglin, H.; Vöcking, B.: On the Impact of Combinatorial Structure on Congestion Games. In: *J. ACM* 55 (Dec. 2008), no. 6.

[AS03] Abdulkadiroğlu, A.; Sönmez, T.: School choice: A mechanism design approach. In: *The American Economic Review* 93 (2003), no. 3, pp. 729–747

[AS08] Ackermann, H.; Skopalik, A.: Complexity of Pure Nash Equilibria in Player-Specific Network Congestion Games. In: *Internet Mathematics* 5 (2008), no. 4, pp. 323–342

[BBD+18] Balogh, J.; Békési, J.; Dósa, G.; Epstein, L.; Levin, A.: A New and Improved Algorithm for Online Bin Packing. In: *26th Annual European Symposium on Algorithms (ESA 2018)*. Ed. by Azar, Y.; Bast, H.; Herman, G. Vol. 112. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018, 5:1–5:14.

[BG21] Bertsekas, D.; Gallager, R.: *Data Networks: Second Edition*. Athena Scientific, 2021.

[BHK17] Brosig-Koch, J.; Hehenkamp, B.; Kokot, J.: The effects of competition on medical service provision. In: *Health Economics* 26 (2017), no. 53, pp. 6–20

[BHK23] Brosig-Koch, J.; Hehenkamp, B.; Kokot, J.: Who benefits from quality competition in health care? A theory and a laboratory experiment on the relevance of patient characteristics. In: *Health Economics* (2023)

[BM14] Biró, P.; McDermid, E.: Matching with sizes (or scheduling with processing set restrictions). In: *Discrete Applied Mathematics* 164 (2014), pp. 61–67

[CDR06] Cole, R.; Dodis, Y.; Rough2006garden, T.: Bottleneck Links, Variable Demand, and the Tragedy of the Commons. In: vol. 60. Jan. 2006, pp. 668–677

[CFGS11] Caragiannis, I.; Fanelli, A.; Gravin, N.; Skopalik, A.: Efficient Computation of Approximate Pure Nash Equilibria in Congestion Games. In: *CoRR* abs/1104.2690 (Apr. 2011)

[CFGS15] Caragiannis, I.; Fanelli, A.; Gravin, N.; Skopalik, A.: Approximate Pure Nash Equilibria in Weighted Congestion Games: Existence, Efficient Computation, and Structure. In: *Transactions on Economics and Computation* 3 (2015), no. 1

[CLS16] Chen, Y.; Liang, Y.; Sönmez, T.: School choice under complete information: An experimental study. In: *The Journal of Mechanism and Institution Design* 1 (2016), no. 1, pp. 45–82

[CS06] Chen, Y.; Sönmez, T.: School Choice: an experimental study. In: *Journal of Economic theory* 127 (2006), no. 1, pp. 202–231

[CS11] Chien, S.; Sinclair, A.: Convergence to approximate Nash equilibria in congestion games. In: *Games and Economic Behavior* 71 (2011), no. 2, pp. 315–327.

[Deb52] Debreu, G.: A social equilibrium existence theorem. In: *Proceedings of the National Academy of Sciences* 38 (1952), no. 10, pp. 886–893

[Del05] Dellarocas, C.: Reputation Mechanism Design in Online Trading Environments with Pure Moral Hazard. In: *Information Systems Research* 16 (2005), no. 2, pp. 209–230

[DRS14] Drees, M.; Riechers, S.; Skopalik, A.: Budget-Restricted Utility Games with Ordered Strategic Decisions. In: *Algorithmic Game Theory - 7th International Symposium, SAGT 2014, Haifa, Israel, September 30 - October 2, 2014. Proceedings*. Ed. by Lavi, R. Vol. 8768. Lecture Notes in Computer Science. Springer, 2014, pp. 110–121.

[DS08] Dunkel, J.; Schulz, A. S.: On the Complexity of Pure-Strategy Nash Equilibria in Congestion and Local-Effect Games. In: *Mathematics of Operations Research* 33 (2008), no. 4, pp. 851–868.

[EHH22] Endres, A. E.; Hehenkamp, B.; Heinzel, J.: *The Impact of Product Differentiation on Retail Bundling in a Vertical Market*. Tech. rep. Paderborn University, 2022

[FFG+18] Feldkord, B.; Feldotto, M.; Gupta, A.; Guruganesh, G.; Kumar, A.; Riechers, S.; Wajc, D.: Fully-Dynamic Bin Packing with Little Repacking. In: *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*. Ed. by Chatzigiannakis, I.; Kaklamanis, C.; Marx, D.; Sannella, D. Vol. 107. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018, 51:1–51:24.

[FHSS18] Feldotto, M.; Haake, C.-J.; Skopalik, A.; Stroh-Maraun, N.: Disaggregating User Evaluations Using the Shapley Value. In: *Proceedings of the 13th Workshop on Economics of Networks, Systems and Computation (NetEcon 2018)*. Irvine, California, USA, 2018, 5:1–5:6

[FKS05] Fotakis, D.; Kontogiannis, S.; Spirakis, P.: Selfish unsplittable flows. In: *Theoretical Computer Science* 348 (2005), no. 2. Automata, Languages and Programming: Algorithms and Complexity (ICALP-A 2004), pp. 226–239.

[FPT04] Fabrikant, A.; Papadimitriou, C.; Talwar, K.: The Complexity of Pure Nash Equilibria. In: *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing*. STOC '04. Chicago, IL, USA: Association for Computing Machinery, 2004, pp. 604–612.

[GHL19] Gu, Y.; Hehenkamp, B.; Leininger, W.: Evolutionary equilibrium in contests with stochastic participation: Entry, effort and overdissipation. In: *Journal of Economic Behavior and Organization* (2019), pp. 469–485

[GMV05] Goemans, M.; Mirrokni, V.; Vetta, A.: Sink equilibria and convergence. In: *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*. 2005, pp. 142–151

[GS62] Gale, D.; Shapley, L. S.: College admissions and the stability of marriage. In: *The American Mathematical Monthly* 69 (1962), no. 1, pp. 9–15

[HHSS14] Harks, T.; Höfer, M.; Schewior, K.; Skopalik, A.: Routing Games with Progressive Filling. In: *Proceedings of the 33rd Annual IEEE International Conference on Computer Communications (INFOCOM'14)*. 2014, pp. 352–360

[HK12] Harks, T.; Klimm, M.: On the Existence of Pure Nash Equilibria in Weighted Congestion Games. In: *Mathematics of Operations Research* 37 (2012), no. 3, pp. 419–436.

[HK20] Hehenkamp, B.; Kaarbøe, O. M.: Location Choice and Quality Competition in Mixed Hospital Markets. In: *Journal of Economic Behavior and Organization* 177 (2020), pp. 641–660

[HK23] Hehenkamp, B.; Kaarbøe, O. M.: *Price Regulation, Quality Competition and Location Choice with Costly Relocation*. Tech. rep. Paderborn University, 2023

[HKS14] Hansknecht, C.; Klimm, M.; Skopalik, A.: Approximate pure Nash equilibria in weighted congestion games. In: *Proceedings of the 17th. International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*. 2014, pp. 242–257

[HLP04] Hehenkamp, B.; Leininger, W.; Possajennikov, A.: Evolutionary equilibrium in Tullock contests: spite and overdissipation. In: *European Journal of Political Economy* 20 (2004), no. 4, pp. 1045–1057.

[Hoo20] Hoof, S.: On a class of linear-state differential games with subgame individually rational and time consistent bargaining solutions. In: *Journal of Mechanism and Institution Design* 5 (2020), p. 1

[HR18] Haake, C.-J.; Recker, S.: The Generalized Nash Bargaining Solution for Transfer Price Negotiations under Incomplete Information. In: *Group Decision and Negotiation* 27 (2018), no. 6, pp. 905–932

[HS18] Haake, C.-J.; Stroh-Maraun, N.: Outcome equivalence in school choice with reciprocal preferences. In: *Economics Letters* 170 (2018), pp. 39–41

[HS20a] Hoyer, B.; Stroh-Maraun, N.: *Matching strategies of heterogeneous agents under incomplete information in a university clearinghouse*. Tech. rep. 2020, pp. 453–481

[HS20b] Hoyer, B.; Stroh-Maraun, N.: *Stability in Weighted College Admissions Problems*. Working Papers Dissertations 63. Paderborn University, Faculty of Business Administration and Economics, May 2020.

[HT21] Haake, C.-J.; Trockel, W.: Socio-legal systems and implementation of the Nash solution in Debreu–Hurwicz equilibrium. In: *Review of Economic Design* (2021)

[Hur94] Hurwicz, L.: Economic design, adjustment processes, mechanisms, and institutions. In: *Economic Design* 1 (1994), no. 1, pp. 1–14

[Ivk95] Ivković, Z.: *Fully dynamic approximation algorithms*. University of Delaware, 1995

[KY18] Koray, S.; Yildiz, K.: Implementation via rights structures. In: *Journal of Economic Theory* 176 (2018), pp. 479–502

[MFHR18] Mir Djawadi, B.; Fahr, R.; Haake, C.-J.; Recker, S.: Maintaining vs. Milking Good Reputation when Customer Feedback is Inaccurate. In: *PLoS ONE* 13 (2018), no. 11

[MM10] McDermid, E. J.; Manlove, D. F.: Keeping partners together: algorithmic results for the hospitals/residents problem with couples. In: *Journal of Combinatorial Optimization* 19 (2010), no. 3, pp. 279–303

[PS07] Panagopoulou, P. N.; Spirakis, P. G.: Algorithms for Pure Nash Equilibria in Weighted Congestion Games. In: *ACM J. Exp. Algorithmics* 11 (Feb. 2007), 2.7–es.

[PS08] Pathak, P. A.; Sönmez, T.: Leveling the playing field: Sincere and sophisticated players in the Boston mechanism. In: *The American Economic Review* 98 (2008), no. 4, pp. 1636–1652

[Rot89] Roth, A. E.: Two-sided matching with incomplete information about others' preferences. In: *Games and Economic Behavior* 1 (1989), no. 2, pp. 191–209

[Sei02] Seiden, S. S.: On the Online Bin Packing Problem. In: *J. ACM* 49 (Sept. 2002), no. 5, pp. 640–671.

[SS74] Shapley, L.; Scarf, H.: On cores and indivisibility. In: *Journal of Mathematical Economics* 1 (1974), no. 1, pp. 23–37

[Str20] Stroh-Maraun, N.: *Pareto Efficiency in Weighted School Choice Problems*. Working Papers Dissertations 64. Paderborn University, Faculty of Business Administration and Economics, May 2020.

[Vet02] Vetta, A.: Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In: *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.* 2002, pp. 416–425

[Wag94] Wagenhofer, A.: Transfer pricing under asymmetric information: An evaluation of alternative methods. In: *European Accounting Review* 3 (1994), no. 1, pp. 71–103

[WLLD05] Wang, J.; Li, L.; Low, S.; Doyle, J.: Cross-layer optimization in TCP/IP networks. In: *IEEE/ACM Transactions on Networking* 13 (2005), no. 3, pp. 582–595

[YXF+10] Yang, D.; Xue, G.; Fang, X.; Misra, S.; Zhang, J.: Routing in max-min fair networks: A game theoretic approach. In: Oct. 2010, pp. 1–10

[YXF+13] Yang, D.; Xue, G.; Fang, X.; Misra, S.; Zhang, J.: A Game-Theoretic Approach to Stable Routing in Max-Min Fair Networks. In: *IEEE/ACM Transactions on Networking* 21 (2013), no. 6, pp. 1947–1959