

FROM “TAKE-IT-OR-LEAVE-IT” OFFERS TO “TAKE-IT-OR BE-LEFT-OUT” ULTIMATUM

A Trade mechanisms for online services

Bodenstein, Christian, University of Freiburg, Kollegiengebäude II, Platz der alten Synagoge, 79085 Freiburg, Germany, christian.bodenstein@is.uni-freiburg.de

Schryen, Guido, University of Freiburg, Kollegiengebäude II, Platz der alten Synagoge, 79085 Freiburg, Germany, schryen@gmx.net

Neumann, Dirk, University of Freiburg, Kollegiengebäude II, Platz der alten Synagoge, 79085 Freiburg, Germany, dirk.neumann@is.uni-freiburg.de

Abstract

In a world, where more and more businesses seem to trade in an online market, the supply of online services to supply the ever-growing demand could quickly reach its capacity limits. Online service providers may find themselves maxed out at peak operation levels during high-traffic timeslots but too little demand during low-traffic timeslots, although the latter is becoming less frequent. At this point not only deciding which user is allocated what level of service becomes essential, but also the magnitude of the service provided, can be controlled by pricing. Pricing is an important factor when efficient and acceptable allocation of resources between individuals must be reached. Without prices, transferring or sharing goods would be impossible. In sharing information, pricing a product however is not as simple as relatively pricing an apple or a pear. Often the costs, and hence the prices are simply unknown. Backed by this scenario, the online services market could be combined with the market design mechanism of diamonds. For this we propose an ultimatum pricing strategy which effectively allows for valuations to be accounted for, but no longer a necessity when pricing in grid, cloud or other online computing environments.

Keywords: Posted Price, Ultimatum Game, Energy Efficiency, Mechanism Design.

1 INTRODUCTION

In a world, where more and more businesses seem to trade in an online market, the supply of online services to supply the ever-growing demand could quickly reach its capacity limits, but only for short time intervals. Online service providers, like Salesforce.com¹, StrikeIron², and eBay³ may find themselves maxed out at peak operation levels during high-traffic timeslots, but too little demand during low-traffic timeslots, although the latter is becoming less frequent. In some cases it might be possible to run certain tasks at night, but for some this is not an option, since some tasks need to be processed in a timely manner. At this point not only deciding which user is allocated what level of service becomes essential, but also the magnitude of the service provided, which can be defined by service level agreements (SLA). This allows for a rather simple method of reducing traffic, by only agreeing on a low service level, - but for some tasks, this is simply not an option. Log records may show that scientific networks peak in traffic before the submission deadlines of journals or workshops, since many scientists attempt to model or test their results before final submittal. Here, a low level of service is, again, of no use. The market can almost be branded as a high-demand type which is in dire need of some sort of energy supply regulation. An estimated 50%-70% (platform dependent) of total 'manageable' costs of datacenters are energy costs. Additively, Gartner estimates, that the combined CO₂ emissions, as a byproduct of energy expenditure, by the ICT-Industry has reached the emissions of the worldwide combined civilian airliner fleet.

Similarly faced with a comparable high-demand, regulated-supply scenario is the market for raw uncut diamonds. The market for trading raw diamonds is held mostly by DeBeers distribution arm, the Diamond Trading Company (DTC), which sorts, values and sells approximately 90 % of the world's rough diamonds by value (Viljoen 2008). The DTC buys the stones from the producers, sells the gem stones to carefully inspected buyers by presenting them with a box of selected diamonds, naming a price, and giving the buyer a take-it-or-leave option. Most buyers' suspect, those who do not buy at a sight are not asked back, effectively dropping out of the market (c.f. TIME MAGAZINE, "Tightest of all" 1945). By buying most of the diamond mines, DeBeers Group effectively created a cartel, in which it can set the prices the buyers have to pay. Notable, is the extremely high demand for diamonds, and an artificially set, low supply. This way the DTC controls the amount of diamonds in the market, and which jeweler is permitted to cut the raw diamonds.

Learning from this, the online services market could be combined with the market design mechanism that is borrowed from trading raw diamonds, as this mechanism seems adequate for online services as well. Transferring this market design to the market for online services, by selecting a specific pricing strategy, the suppliers could be reduced to only efficient suppliers, much like the DTC chooses only the best cutters. Combined with a prioritization of online service demands, merged with a green scheduling principle the energy requirements by the sector could be reduced, without forfeiting too many economic incentives and properties.

The research question posed in this paper is: "*Can 'take-it-or-be-left-out' (short TIBLO) trade mechanisms be applied to the market of online services to achieve favorable results?*"

Let the diamond miners be replaced with nodes supplying online services, the sight holders substituted with agents requiring services and DeBeers Group is exchanged with a scheduling agent (automated or as a virtual third party). Potential in this trading mechanism is the strong control the scheduler has over the price and output of the online services, similar to the control the cartel has over the diamond price and gem supply. To this point, the game is identical to a standard posted-price auction. New to the game is the threat to expel the agent which does not accept the offer posted by the scheduler. As will

¹ <http://www.salesforce.com/> - CRM Software on Demand

² <http://www.strikeiron.com/> - Data as a Service

³ <http://www.ebay.com/> - New & used electronics, cars, apparel, collectibles, sporting goods & more at low prices

be shown in this paper, this threat is what keeps both agent parties from bidding/asking under-/overvalued prices, resulting in fair prices.

In previous work we showed a model which offered an energy efficient approach to scheduling in grid/cloud environments (self-citing). We assumed that all agents know their valuation of a single unit of CPU and were at least willing to post this offer. As a derivation of the greedy heuristic, the bidder with the highest valuation was allocated to the node with the least energy costs, followed by the second highest bidder who was allocated to the second most efficient node, etc. Table 1 shows an extract of the example schedule using the green heuristic algorithm, where each job is of unit size, and each job has 2 processing units available. For simplicity we assume a finite single period model:

<i>Nodes</i>			<i>Jobs</i>		<i>Job Allocation</i>		
<i>Node</i>	r_n	ϵ_n	<i>Job</i>	b_j	<i>Node</i>	<i>Job</i>	<i>price</i>
N1	9	8	J1	20	N1	J2	17
N2	11	4	J2	17		J4	13
N3	10	12	J3	11	N2	J5	24
			J4	13		J1	20
			J5	24	N3	OFFLINE	

Table 1. Sample Job allocation using the green heuristic algorithm

Applying the greedy heuristic (Stöber et al. 2007), would result in an allocation N1(J5, J1), N2(J3) and N3(J2, J4), while yielding a higher revenue, has considerably higher energy costs than the green heuristic, where the inefficient node was shut down, and J3 was not allocated since it was too expensive. We found however that not all agents had an incentive to truthfully report their valuations, let alone the fact that some agents may not even be able to value CPU power on their own since power usage by electronic devices, without the use of special measuring devices are rather concealed.

In this work, we wish to propose an alternative trade mechanism to counter this incentive to misreport and unknown valuations, by proposing an ultimatum pricing strategy which effectively allows for valuations to be accounted for, but no longer a necessity when pricing in grid, cloud or other online computing environments. We will restrict our analysis only to the impact of posted price trades coupled with an ultimatum condition. Hence the contribution of this paper is threefold:

- Devise a **simplistic model** that captures the essential problems in bidding strategies without being too complex to be tractable
- Revise current pricing models to incorporate **energy efficiency** in their objectives
- Incorporate posted price mechanisms to allow buyers with **unknown valuations** to participate, while keeping the bidding mechanism to allow for **urgency signalling**

The paper is structured as follows. After a brief introduction to the topic, section 2 contains the requirements to online services and motivational scenarios, where the model may be applicable. In section 3, we will briefly discuss auctions and posted price mechanisms, and related work in this field. In section 4 we present a game-theoretic model showing the strategic implications of a posted price ultimatum game. The evaluation based on the motivational scenarios presented is covered in section 5 followed by concluding remarks in section 6 which wrap up this work.

2 MOTIVATIONAL SCENARIOS AND REQUIREMENTS

Pricing is an important factor when efficient and acceptable allocation of resources between individuals must be reached. Without prices, transferring or sharing goods would be impossible. In

sharing information, pricing a product however is not as simple like relatively pricing an apple to a pear. Goods can be weighed, used, thrown, - all properties which are of definite value to somebody. Information goods, or virtual services on the other hand are harder to price. In some cases users are not even aware of the costs involved, like for example the costs involved to perform a simple Google search are still largely unknown to the general user. The energy costs, of performing a Google search, are the equivalent of powering a light bulb for an hour. Based on this information on energy costs to transfer, store, update or use information in online services, pricing strategies can be derived, by relating them to the energy costs of the service. A positive side-effect of this approach is its immediate optimization of profit, directly dependant on the energy costs. This approach can be used for a number of scenarios in online services:

2.1 Motivational Scenarios

This paper is structured following the Design Science guidelines (Hevner et al. 2004) in attempt to develop a viable model and methodology for solving the important and relevant problem of high energy costs throughout the ICT industry. It is therefore necessary to show a few scenarios which firstly show what is meant by the term online services, and secondly already provide a rough idea of what is required of a model to function in the market for online services. To preserve generality, we found our analysis on the *pricing channels*, *cloud computing*, *Mashups* and *hosting services*.

2.1.1 “Name Your Own Price” Channel

“Name Your Own Price®”⁴ is a pricing mechanism where interested buyers pick a commodity and post an offer of how much they are willing to pay for that good. The seller may then accept or reject the offer. In practice, this mechanism results in slightly lower than market value prices, but not by much. Rational buyers therefore must obviously value the goods fairly. Priceline.com and its famous ‘Name Your Own Price®’ service is one of the leading travel service providers for price-conscious travelers, which claims to offer the cheapest possible travel arrangements. Users input their desired departure location, destination, time horizon and price they wish to pay. They are then offered the cheapest travel package. Problematic however is the fact that the bids are not binding, in the sense that potential buyers use these portals more as an information source rather than buying directly. In this scenario, posted-price mechanisms coupled with an ultimatum shown in this paper could be useful, to bind customers to the service, while ensuring a fair match of low price offers through reputation.

2.1.2 Cloud Computing

The term “Cloud Computing” originated from diagrams where the technology architecture depicted the Internet as a “cloud” of services, describing how most users do not really see the hardware they access, therefore often being perceived as working “in the clouds”. Consumers of cloud computing services purchase computing capacity on-demand and are not concerned with the underlying technologies used to achieve the increase in server capability; they purchase the software as a service (SaaS). Current cloud computing providers like Amazon.com⁵ or Sun offer computing resources at fixed prices and are contracted by SLA’s, but unified pricing strategies have yet to be implemented. Posted-price auctions could offer an efficient pricing strategy which promotes efficiency and optimality when coupled with the right scheduling design. With respect to energy efficiency, the posted-price auction coupled with the green scheduling algorithm could prove useful and advantageous in achieving a much needed reduction of energy costs of data centers.

⁴ ‘Name your own price’ mechanisms are registered service marks of priceline.com Incorporated.

⁵ <http://aws.amazon.com/s3/> - Amazon Simple Storage Service (Amazon S3)

2.1.3 *Mashup-as-a-Service*

A Mashup is a web application which combines data from multiple sources onto one interface. An example would be if a real-estate owner would link his location information onto Google Maps and providing it as a service. As the term Mashup implies it is an easy and fast, “mashed” integration of numerous data sources to provide a service neither of the original data sources intended. The concept of Mashup-as-a-Service, similar to cloud computing, would entail producing on-demand services, as the users need them, using third-party sources as inputs. A good example of on-demand pricing of services is Serena.com⁶, which provides business users with Web 2.0 tools to build business Mashups. Pricing is currently done through a subscription on a “per User, per Month” basis. Discriminatory pricing could attract more small scale users, who wish to simply provide for a ‘fire-and-forget’ application. Other examples could include Yahoo Pipes⁷, which is currently free of charge.

2.1.4 *Hosting Services*

Hosting services generally refer to web hosting services where many websites are hosted on one web server connected to the Internet. This is generally the most economical option for hosting as many people share the overall cost of server maintenance. Although posted-price strategies have been extensively used in modern markets, including an ultimatum could prove useful in pursuit of efficient allocation of resources.

2.2 **Requirements for Selling Online Services**

With the motivational scenarios defined, we can derive the requirements for online services. These requirements are set by the motivational scenarios and are real problems faced by the online services market (some contribute more to the importance of requirements than others), and are the absolute minimal requirement set to the model presented in this work. The mechanism is thus required to be compatible with scenarios of:

1. **Excess Demand / Supply** (Sub-Market dependent). Resource allocation is not a problem, if sufficient resources are available. Especially for hosting services and Clouds, this is a big issue. The requirement is thus: “Are the strategies imposed under “normal circumstances” still dominant in cases of excess demand / supply?”
2. **Single bids**. Contrary to requirement 1, the markets are not always saturated in demand/supply. Often they are underrepresented: “Does the mechanism still function when faced with only one buyer/seller against many sellers/buyers?”

Obviously requirements 1 and 2 reflect the highly volatile nature of the market for online services. While an internet service provider finds that his traffic peaks around mid-day, when all employees check their personal e-mails, he might find that his traffic is lowest at 3 o’clock in the morning when most customers are asleep. Further, the model must meet the requirements of:

3. **Energy efficiency**. Foremost, besides having feasible properties for market trading, the mechanism must promote the use of energy efficient allocation mechanisms.
4. **Individual Rationality**. An important requirement for the mechanism is that both trading parties benefit from trade, and rational behavior is induced. By assumption, all agents are taken as risk-averse.
5. **Unknown valuations**. It is important that the mechanism can process jobs with unknown or withheld valuations, since some buyers may be unaware of their own valuations. This could prove especially useful in the case of Mashups-as-a-service, where some tasks are so unique,

⁶ <http://www.serena.com/mashups/index.html> - Serena Business Mashups

⁷ <http://pipes.yahoo.com/pipes/> - Rewire the web

that pricing cannot be done given observation of historic prices, but also in other aspects of trading online services.

6. **Complexity.** The mechanism should not introduce any unnecessary complexity, nor add any intractable problems to the trading environment. This is a must for all scenarios! It is of no use to derive a complicated trading scheme which is of no use because no-one can use it.
7. **Time preference.** The users require resources in a timely manner; hence the trade mechanism may not take too long to decide which resource is allocated to a task, to avoid encroaching on the task deadline.

If the above requirements are met, the TIBLO pricing mechanism presented in this work could prove a favorable alternative to current pricing mechanism used in trading online services. In the next section, we will briefly discuss auctions and posted price mechanisms, and related work.

3 RELATED WORK

Current research in Grid Computing has come up with numerous scheduling mechanisms to efficiently schedule supply and demand of computing resources. To provide incentives for resource suppliers, pricing mechanisms have been proposed and seem the most promising method to provide incentives for sharing resources. Although some networks (Folding@home, SETI@home etc.) do not compensate in payments, they are still able to attract resources by other means (interest, appeal to moral duty, etc.)

For commercial use however these incentives are insufficient and often infeasible since a rational supplier would not willingly and knowingly forfeit profits and freely give them to a third party scheduler, who in turn makes money with it. Auctions, which have previously not only been used, to transfer assets from public to private hands, but also to uncover the valuations of bidders, seemed most promising in this regard. As long as each bidder knows his own valuations in form of a private value, the outcome of the auction in most cases awards the good to the bidder with the highest valuation (Krishna 2002). Without this information, sellers can only guess the average willingness to pay for each good they wish to sell, - by posting a fixed price.

Wang (1998), among others, compared these two popular selling methods (auction and posted price) in a one-period correlated valuation model, and found that auctions were mostly preferred, since static pricing has been labeled as economically inefficient, as it does not include pricing the dynamic nature of resource requirements. In this work however we adapt static posted price to the discriminatory posted price auction mechanism (Seifert, 2006). Several authors have investigated hybrid market institutions which combine auctions and posted price offers. Combining auctions with a posted buyout option was first noted by Lucking-Reiley (2000). Hinz et al. (2008) developed an analytical model analyzing the effect of shared information on bidding behavior in a price auction variation of Name-Your-Own-Price addressing the often hidden valuation still inherent to standard bids.

According to Wang et al. (2004) combined auctions are not a new phenomenon, since in a pure auction institution the posted price can be set so high that it is never even considered as a bid, and in a posted price institution it is set so low, that no bids exist below it, implying that pure posted price or auction institutions are merely a polarization of the combined mechanism.

These models however all rely on the assumption that the posted price depicts the highest possible achievable price, and it makes no sense to bid beyond the posted price (Seifert, 2006) since buyers prefer the posted price over the auction mechanism, if the posted price is less (or equal) to their own valuation of the product with at least a weak dominance.

Currently there is no approach that satisfies all requirements imposed in section 2. In the next section we will present our own price-setting mechanism, which includes all the favorable properties of the APPO (Seifert, 2006), but is simple enough to be used in a fast and volatile market of trading online services, while still concurring to the ideals of 'green' scheduling algorithms, instead of profit maximization strategies.

4 THE MODEL

The model discussed, shows an alternative pricing mechanism for online services. Using Amazon.com as an example, where the “task request” to be allocated is the storage of data from research, and the “suppliers” is the host of a server farm, the mechanism can be described as follows: A scientist wishes to store his research on a server at Amazon.com. She approaches Amazon and is faced with two options: Reveal her valuation, thereby influencing the probability her job is accepted, or simply accept the posted price by Amazon, which is subject to the prices of jobs currently on the system. Once the price has been determined, she can either accept the package offer, of storing her data or reject. Upon rejection, her personal data is stored only such that Amazon can disclose her from ever trading at Amazon.com in future. This example simply explains the model we wish to present more formally in the following.

In section 4.1 we initiate the benchmark model more formally and strategically, where only the strategies of a single negotiation between a unit buyer and seller are uncovered. Section 4.2 discusses the effect of adding outside options to the benchmark model, followed by section 4.3 which includes analysis of scenarios with uncertainty.

4.1 Benchmark Model

In a benchmark bilateral monopoly setting, which corresponds to a market with one buyer (B) and one seller (A), the solution is rather simple. The individual, both as a consumer and producer is a small part of the market and can therefore take everyone else’s behavior as given; the world is merely a set of prices (Friedman 1990). The actions of the two individuals could therefore be depicted in a game. To describe the posted price, ultimatum game, an extension to a classical ultimatum game, we use a normal-form representation of a two-player sequential game Γ after a pricing decision by A , where each player (Player 1 $\equiv A$, Player 2 $\equiv B$) can choose between two strategies ($S^A = \{s_1^A, s_2^A\}$, $S^B = \{s_1^B, s_2^B\}$) once B has issued a request. The game is set up as follows:

- Player B sends a job request to A with job specifications and some form of signal of valuation.
- Player A offers a price $p \in [0; P]$.
- After A has posted a price, B can either choose to accept (a) the offer, or reject it (r).
- Subject to S^B , A can choose to trade with B again (a) or expel him from trade (r).

In other words the two strategies represent the agent’s choice between accepting the proposal or not.




Figure 2. Game tree of the standard posted-price game.

The strategy space S is defined by the Cartesian product of the individual strategies of the two players:

$$S = S^A \cdot S^B = \{(a,a), (a,r), (r,a), (r,r)\}$$

By game design, if Player A has the world’s only online service available and Player B can be summarized into an aggregate buyer, the outcome is simple. The online service is worth nothing to A , but is worth $p = v_B$ to B , resulting that when both reject, nothing is won: In a one-shot game, both players will accept, since there is no outside option available, and without trade both are worse off as long as the posted price $U_B(p) \leq U_B(v_B)$. In this setting for $p \geq 0$ any rational buyer will always accept.

Unique to this type of ultimatum game, is the convergence of welfare distribution to the 50/50 Schelling point (Schelling, 1963), similar to the pie distribution game. Consider A has energy costs of executing a unit size job of 1 €, which is of value to B up to 3 € and the information is known by both agents. Monopoly pricing would suggest A charges a price of $p = 3$ (full welfare benefits to A) while competitive pricing would suggest $p = 1$ (full welfare benefits to B). Fair pricing would suggest equal distribution of welfare, resulting convergence towards $p = 2$, with at least a weak dominance. A however still does have some form of power, and may choose to charge slightly higher prices than $p = 2$, but as $2 < p \leq 3$ the probability that B rejects the offer increases.

Theorem 1: *The threat to expel bidders from the game if they do not accept is credible.*

Proof sketch: Proving whether or not the threat of expulsion is credible or not is rather simple, following Schelling's (1963) strategic principles of *continuous negotiations* and *casuistry*. The former is solidified by the public knowledge that any buyer who refuses the offer is expelled. According to Schelling (1963), the party threatening achieves commitment to execution not by the gains from actually executing the threat, but by pointing out the long-run value of accepting, regardless of his own losses. Hence for all $s'_i(.) \in S$, $S(a, a)$ is a dominant strategy Nash equilibrium since $u(s_i(a), s_{-i}(.)) \geq u(s'_i(.), s_{-i}(.))$. ■

Theorem 2: *Even in his monopoly position the scheduler has a weakly dominant strategy to charge "fair" prices.*

Proof sketch: *Casuistry*, as a form of rationalization (Schelling, 1963), ensures that agents will even attend the market, in that the scheduler denies himself too great a reward from his monopoly position (which in effect is bolstered by the take-it-or-leave-it mechanism), effectively lowering his prices to a reasonable level: at the least, lower than profit maximizing prices. This becomes an important pricing principle when the scheduling agent is no longer the sole supplier of computing power. ■

More formally, to maximize $(b_j - e_n)$ the highest bidder should be allocated the most efficient node, which effectively results in an optimization of energy costs in the system, the pricing decision can be denoted mathematically by the following optimization function,

$$\begin{aligned} \max_{x, p} V &:= \sum_j c_j \sum_t \sum_n x_{jnt} [p_{jn}(b_j, e_n)] = \sum_j c_j \sum_t \sum_n x_{jnt} \left(\frac{b_j - e_n}{2} \right) \\ \text{s.t.} \quad &x_{jnt} \in \{0, 1\}, \quad \sum_n x_{jnt} \leq 1 \quad \forall j \in J, t \in T \end{aligned}$$

where c_j denotes the amount of CPU resources required, x_{jnt} corresponds to the decision variable unique for each time slot t , p_{jn} is the discriminatory price subject to b_j and e_n , b_j represents the bid by agent j and e_n stands for the energy costs per CPU of node n . Additionally, $b_j \geq r_n$. The atomicity, and time-feasibility constraints have been omitted, since for this case we look only at unit-size jobs which all have a duration of one time unit.

4.2 Outside options

The simple bilateral monopoly is now introduced with a third party service provider, like Amazon Web Services for example, which sells all services at a fixed price p_s and pursues no further strategies but to accept all jobs until its capacity limit k_s is reached. This third party is known to schedule its jobs in a "first-in-first-out" (henceforth referred to as the 'FIFO-agent') way, ignorant of its energy costs or CO₂ emissions (this is not to say that Amazon is ignorant of energy costs). The TIBLO-agent (referred to as the 'Seller' in previous chapters) who includes energy efficient scheduling principles is now faced with the following payoff structure:

A\B	A	R
a	$p, (v_B - p)$	$0, (v_B - p_s)$
r	$0, (v_B - p_s)$	$0, (v_B - p_s)$

Table 3. Posted-price payoff matrix

Hence when posting a price p , to a rational buyer who is ignorant of energy efficient principles, $p \leq p_s$ which would result in an equilibrium strategy at $S = \{(a, a)\}$. Faced with an energy-aware buyer, A has to ensure that $p \leq p_s + \delta$, where δ is the valuation of energy-awareness, which allows the TIBLO-agent to charge slightly higher prices than the FIFO-agent. This phenomenon can be found in many product pricing strategies where environmental factors are included in the product. *Badenova*⁸, a german power distributor for example, charged a higher price for their product “*Regiostrom*” (which draws 50% of its power from regenerative resources) than their standard price power, which draws power from conventional ‘unclean’ sources.

Since the FIFO-agent cannot discern the buyers without drawing some sort of signal prior to the pricing decision, the exploitation possibility of δ are limited. Hence $0 \leq p_s \approx p \leq P$, where setting $p = 0$ means that the buyer receives the full benefit of trade, and P implies full benefit to the sellers.

Shown in its normal form, where the relative utility is $p - p_s$:

A\B	a	R
a	$++ , \pm$	$0, -$
r	$0, -$	$0, -$

$p < p_s$

A\B	a	r
A	$++ , -$	$0, \pm$
R	$0, +$	$0, \pm$

$p > p_s$

Table 4. Posted-price relative utility matrix

Even with outside options, theorem 1 and 2 still hold. Since the TIBLO-agent is faced with an outside option, the credibility of the threat to expel the agent only diminishes with $p \rightarrow P$. As implied by *casuistry* (Schelling, 1963) and reinforced by upholding his reputation to expel rejecting buyers, his pricing strategy is now limited to the space $(0; p_s + \delta)$, where $\delta = 0$ for ignorant buyers, and $\delta > 0$ for energy-aware buyers. Again, since the TIBLO-agent cannot discern ignorant from aware buyers, to keep his threat credible, the TIBLO-agent is forced to abandon his price setting position, resulting $p \leq p_s$. This is an important result since by pledging the threat of expulsion, the TIBLO-agent, given outside options, effectively loses his price setting trait. Knowing this, buyers should continue to approach the TIBLO-agent for online services. In section 4.3 the model presented is extended to include more than one interested buyer with different valuations.

4.3 Bayesian case

The simple bilateral monopoly is now introduced with a third party service provider and at least one further buyer interested in the online service product. Without loss of generality let's assume that the single, continuum of energy-aware buyers is split into two aggregate buyers, B^L representing the low valuation buyers with v_B^L and B^H representing the high valuation buyers with v_B^H , the game outcome is somewhat more complex, where A does not know in advance whether B is of type L or H . In fact, B does not even know his relative valuation position to other bidders, but only his own valuation. To describe this scenario, the game is extended to a Bayesian TIBLO game with incomplete information, which can be depicted as follows:

⁸ <http://www.badenova.de>




Figure 5. Game tree of the extended Bayesian TIBLO game.

More formally, as a game of incomplete information, the TIBLO-game consists of 2 players, similar to the standard posted price game and a probability π which determines the nature of the buyers, transforming the game of incomplete information into one of imperfect information (Harsanyi 1968).

Essentially, the strategies imposed by the benchmark TIBLO model still hold for the Bayesian version of the game. The agent is forced by his own credibility to offer such prices that the trading partners will definitely agree on. Since B maximizes his utility based on only the current information available to him, the sub-game for each B with A is merely the benchmark game, with one seller and one buyer.

From A 's point-of-view however he is faced with two or more buyers for a finite amount of goods. The decision which job is scheduled is derived from the bidders' valuations. For example, A who has a single unit of CPU available, is faced with two buyers B_1 and B_2 who both request resources for a unit-size job, which can be processed with just 1 unit of CPU. The power costs of A are 2€ and B_1 values his job at 3€ and B_2 at 5€. Given full information, A should invite and offer the CPU to B_2 at a price of 3.50€, who should as a rational buyer accept the offer. But given incomplete information, A merely knows the probabilities that the buyer currently under decision is either the one with a high valuation or a low valuation. The dilemma faced by A is as follows:

- Post a price of 2€ which would satisfy both the low and high valuation buyers, and definitely ensure that both players would accept if asked but forfeits a possible surplus of 1.50€ if he is faced with a high valuation buyer.
- Post a price of 3.50€ which would cause B_1 to reject with positive probability, but still be accepted by B_2 . In this case however A forfeits all future trade opportunities with B_1 .

Generally it can be shown that if the TIBLO-agent is faced with two buyers for a single commodity, and expects a necessity of future trades, he has a strong incentive to keep both agents. This is largely due to the high value of future trades, which exceeds the surplus in a single trading opportunity.

For too high prices, the buyers will call the threat, and reject: If at this point the TIBLO-agent does not expel the partners, he loses all credibility and no longer is able to set the prices. This however is only the case if trading is transparent, and all buyers know of the actions of all other buyers in the market. Pending some sort of communications platform, this is however rarely the case. Therefore, each buyer has to act in maximizing his own utility, which results in the same principle stated in theorem 1 and 2.

5 IMPLICATIONS FOR ONLINE SERVICES

In this work we proposed the use “take-it-or-leave-it” trade mechanisms to be applied to the market of online services, combined with the energy efficient algorithm to effectively reduce energy consumption by online services without forfeiting too much return. The implications for online service providers compared to the *status quo*, using the TIBLO-pricing strategies are threefold:

- Possibility and environmental obligation to reduce energy inefficient allocation of resources to operating tasks
- Self-induced, yet still profitable control of monopolies, resulting in fair, welfare maximizing pricing strategies
- A steady, quantifiable, and most importantly, a forecast-able demand, thanks to the binding properties of the ultimatum mechanism allowing for optimal resource availability planning.

Therefore the TIBLO-mechanism has proven to be not only a viable alternative, but a real first-choice option when prioritizing the reduction of high energy costs, as it adhered to all the requirements set out in section 2.2. For the motivational scenarios, this implies the following:

- *“Name your own price” Channel* – By including an ultimatum to the search mechanism, price-scouting behavior by free riders can be reduced, since they will naturally reject the offer by Priceline.com, for example of a flight from Germany to Hawaii, and be banned from the system as a result. This way, name-your-own-price online services ensure that users will offer real prices, which ultimately could lead to an improved trading platform.
- *Cloud Computing* – Clouds and Grids yet have to decide on a common pricing system. A first-mover to employ TIBLO-pricing strategies would have a significant advantage than if he introduced the system in an already embedded market, where the new pricing mechanism may come across skepticism. TIBLO-pricing could offer an efficient pricing strategy which promotes efficiency and optimality when coupled with the right scheduling design resulting in achieving a much needed reduction of energy costs of data centers, as well as economically viable and welfare-distribution maximizing prices.
- *Mashup-as-a-Service* – To best illustrate the benefits of deploying TIBLO-pricing strategies by Mashup providers, we pick up the example of Mashups-on-demand by Serena.com. Pricing is currently done through a subscription on a “per User, per Month” basis. Discriminatory pricing could not only attract more small scale users, who wish to simply provide for a ‘fire-and-forget’ application, but also increase profits relative to decreased energy costs.
- *Hosting Services* – Although the market for hosting services is already embedded by posted prices offers, the addition inherent to the TIBLO-mechanism could prove a very effective tool to reduce energy costs of datacenters while keeping service profits untouched.

By being able to control the allocation, and effectively, the distribution of resources, one might label the mechanism as a centrally planned micro-economy. Political economists often argue that central planning is the only way that rational buyers can be forced to choose environmental optimization over profit maximization, since CO₂ emissions, for example, are externalities which are not included in the optimization calculation. Additionally, environmentally friendly solutions are often more expensive than unclean sources.

Allowing a corporate individual to control the market in such a way, standard microeconomic theory suggests that market power would transform this individual to a monopolist, resulting in monopoly prices. This is not the case with the TIBLO mechanism, as by issuing a threat of expulsion, the monopoly becomes self-regulated, as shown in the above model. The resulting prices and allocation however still remains feasible and profitable.

6 CONCLUDING REMARKS

In this work we presented a mechanism which presents online service providers with a pricing model combined out of a posted price ‘auction’ and an ultimatum. It presents service providers with the option to reduce energy costs of their resource centers by invoking a discriminatory pricing mechanism which allows for a ranking of requests, without immediately confounding to direct monopoly pricing. This is largely due to the effect, that the provider must uphold his credibility of expelling agents who refuse to abide to his terms. In fact, the resulting self-imposed 50/50 ‘Schelling price’ results in a welfare maximum.

Even without a valuating bid, the TIBLO-pricing can still offer a standardized price package, obviously to lesser terms to keep up the incentives to bid. Also by binding the users to the service, suppliers are more able to foresee demand spikes, and prepare accordingly, lessening the impact of excess demand spikes, possibly through prioritization.

Most importantly, the model can achieve its goal of allocating important (signaled by the valuation) jobs to efficient nodes, which constitutes an important design goal for all future IT-services, Green-IT.

Regardless, in future work we intend to expand the current TIBLO model to include further aspects which may be important for online services, or even other markets. This includes:

- Extending the TIBLO mechanism to include a bidding possibility to pay prices larger than the posted price, which would include a further urgency lever based on the valuation of others in direct competition over resources in our model, knowing the bids of others.
- Expanding the application scenarios to other trading environments in the online services market
- Performing an online field study to evaluate the real response of bidding agents who are not necessarily always risk averse or even irrational at times.

References

- Friedman, D.D. (1990). Price Theory: An Intermediate Text. 2nd Edition. South-Western Publishing Co., Cincinnati, Ohio.
- Harsanyi, J.C. (1967-68). Games with Incomplete Information Played by "Bayesian" Players, I-III. Management Science, 14 (3,5,7), 159-182, 320-334, 486-502.
- Hevner, A. March, S. Park, J. and Ram, S. (2004). Design Science in Information Systems Research. MIS Quarterly, Vol. 28 No. 1, pp. 75-105/ March 2004
- Hinz, O. Spann, M. (2008). The Impact of Information Diffusion on Bidding Behavior in Secret Reserve Price Auctions. Information Systems Research Vol. 19, No. 3, September 2008, pp. 351–368
- Krishna, V. (2002). Auction theory. Academic Press, San Diego, CA
- Lucking-Reiley, D (2000): Auctions on the Internet: What's being Auctioned and How?, Journal of Industrial Economics, 48(3), S. 227-252.
- Mas-Colell, A. Whinston, M.M. and Green, J.R. (1995): Microeconomic Theory, Oxford University Press, New York, NY.
- Myerson, R.B. (1978). Refinements of the Nash equilibrium concept. International Journal of Game Theory, 7 (2), 73-80.
- Myerson, R.B. (1981). Optimal Auction Design, Mathematics of Operations Research, 6(1), S. 58-73.
- Schelling, T.C. (1963). The Strategy of Conflict. Harvard University Press, Cambridge, Massachusetts
- Seifert, S. (2006). Posted price offers in internet auction markets. Lecture notes in economics and mathematical systems, 580. Springer Verlag, Berlin.
- Stößer, J. Neumann, D. Weinhardt, C. (2007) Market-Based Pricing in Grids: On Strategic Manipulation and Computational Cost. Journal of AIS, Sponsored Theory Development Workshop TIME, Tightest of all, 12 February 1945, available at <http://www.time.com/time/magazine/article/0,9171,852161,00.html?iid=chix-sphere>.
- Viljoen, C. (2008). [Evolution of Market Power In The Diamond Industry Supply Chain \(800 B.C. – 1930 A.D.\): Why De Beers Was So Different](#). Economic Research South Africa Networking Workshop, 25-26 September 2008, Durban, South Africa.
- Wang, X. Montgomery, A.L. and Srinivasan, K. (2004). When Auction Meets Fixed Price: A Theoretical and Empirical Examination of Buy-it-Now, Discussionpaper, Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, PA, available at <http://www.andrew.cmu.edu/user/alm3/papers/buyitnow.pdf> (November 23, 2008)
- Wang, R. (1998). Auctions versus Posted Price Selling, American Economic Review, 83(4), S. 838-851.