

Mechanism Design for Task Procurement with Flexible Quality of Service

Enrico H. Gerding¹, Kate Larson², Alex Rogers¹, and Nicholas R. Jennings¹

¹ University of Southampton, Southampton, SO17 1BJ, UK
{eg,acr,nrj}@ecs.soton.ac.uk

² University of Waterloo, Waterloo, ON N2L 3G1, Canada
klarson@cs.uwaterloo.ca

Abstract. In this paper, we consider the problem where an agent wishes to complete a single computational task, but lacks the required resources. Instead, it must contract self-interested service providers, who are able to flexibly manipulate the quality of service they deliver, in order to maximise their own utility. We extend an existing model to allow for multiple such service providers to be contracted for the same task, and derive optimal task procurement mechanisms in the setting where the agent has full knowledge of the cost functions of these service providers (considering both simultaneous and sequential procurement). We then extend these results to the incomplete information setting where the agent must elicit cost information from the service providers, and we characterise a family of incentive-compatible and individually-rational mechanisms. We show empirically that sequential procurement always generates greater utility for the agent compared to simultaneous procurement, and that over a range of settings, contracting multiple providers is preferable to contracting just one.

1 Introduction

Service-oriented computing, in which computational resources are seamlessly and dynamically procured from third party suppliers as they are required, has generated significant recent activity within the research community. Examples of such initiatives include Grid and utility computing, and these technologies are increasingly being proposed for a wide range of scientific and business workflows. However, to reach the full potential of this vision, such systems require that both the suppliers and consumers of these computational resources are able to engage in autonomous negotiation and contracting (given their own individual goals and requirements). To this end, agent-based approaches that make use of computational mechanism design have been advocated [1].

Much of the work in this area to date has sought to extend the standard approaches of mechanism design to cases in which there is a non-zero probability that a service provider may fail to meet its contracted obligation (e.g. failing to satisfy an agreed time deadline) [2,3]. However, a significant shortcoming of much of this work is that it assumes the probability with which a service provider will fail is fixed and exogenous. In contrast, Matsubara provides a more realistic model in which self-interested service providers flexibly manipulate their quality of service in order to maximise their own utility [4]. In this model, a contracted service provider actively manages the resources

that it commits to a task, in order to manipulate the quality of service that it provides to the contracting agent (e.g., it may commit more resources to a task, increasing the probability that the task will be successfully completed, if the reward for doing so is large; conversely, it may intentionally fail to complete a task if it profits by doing so). Here the contracting agent is faced with a principal-agent problem since it cannot directly monitor the service providers' activities. Hence it must create appropriate incentives in order to achieve the desired quality of service level.

To this end, Matsubara provides a mechanism based on a payment rule that incentivises potential contractors to truthfully reveal their costs for completing the task, and, once selected, to invest the required amount of costly resources. The approach used is similar to that more recently applied to an information setting where *strictly proper scoring rules* are used to incentivise information providers to truthfully reveal a probabilistic estimate whose generation requires the investment of costly resources [5,6]. However, these mechanisms are restricted to the case that a single service provider is contracted to complete each task. In reality, when faced with the uncertain execution of tasks, it is common to introduce redundancy by contracting multiple providers for the same task; either simultaneously or sequentially (i.e. the agent awaits the failure of one provider before approaching another) [7].

It is this shortcoming that we address in this paper, and to this end, we describe a novel family of mechanisms that allow the contracting agent to procure computational tasks from multiple self-interested service providers. We consider how a utility-maximising contracting agent should price contracts with these service providers in both the full information setting, where the agent has complete knowledge of the cost functions of the providers, and also in the incomplete information setting where it must elicit this information from them. Such an extension is challenging since, in the optimal case, the expected utility of a provider that has been contacted to perform the task, depends on the cost functions of the other contracted providers. In the incomplete information case this results in an interdependent valuation setting with so-called allocative externalities, for which it has been shown that no standard mechanism exists which is both efficient and incentive compatible [8]. We address this by developing mechanisms which are not efficient, but take advantage of the optimal solution to reduce this inefficiency. We then show empirically that, under our mechanisms, contracting multiple service providers for the same task is preferable to contracting a single provider, as it often increases the probability that a task is successfully completed, whilst reducing the costs to the contracting agent.

In more detail, we make the following contributions to the state of the art:

- We derive optimal task procurement mechanisms when the agent has full knowledge of the cost functions of the service providers. We consider procurement from multiple providers, and consider settings where these providers are contracted simultaneously and sequentially (i.e. the agent awaits the failure of one provider before approaching another).
- We extend these results to the incomplete information setting where the agent must elicit the cost information from the providers. We characterise a family of mechanisms and prove that they are incentive compatible (i.e. the providers have a dominant strategy to truthfully reveal their cost functions to the agent), and individually

rational (i.e. the expected utility of providers that participate in the mechanism is greater or equal to zero). Based on the insights obtained from the optimal, full information case, we present three mechanisms from this family: a uniform and discriminatory pricing mechanism for the simultaneous procurement case, and a mechanism for the sequential procurement case.

- Finally, we empirically evaluate our mechanisms, showing that sequential procurement always generates greater utility for the agent compared to simultaneous procurement, that discriminatory pricing always generates greater utility for the agent than uniform pricing, and that over a range of settings, procuring from multiple providers is preferable to procuring from just one.

The remainder of this paper is structured as follows. In section 2 we formally describe the setting which we consider. In section 3 we describe the optimal procurement strategy in the full information setting, before presenting our three novel mechanisms, and proving their properties, in section 4. We instantiate and empirically evaluate these mechanisms in section 5, and, conclude in section 6.

2 Problem Description

Our model is based on that of Matsubara, but expressed here in the standard notation of mechanism design, rather than the original contract theory. Hence, we consider that contracting agent, A , has a task, T , that it would like to have completed. If the task is completed successfully then the agent receives value V , and otherwise it receives zero. We assume that there are n service providers capable of performing task, T . The probability of any provider successfully completing the task (the quality of service offered) depends on the amount of some costly resources that it decides to allocate to the task.

Formally, we assume that each service provider, i , has a potentially unlimited supply of resources¹, and denote $r_i \geq 0$ as the amount of resources that i will devote to executing the task. We assume that as a provider allocates more resources to the task, the probability that the task will be successfully executed increases. That is, there is a *quality of service* (QoS) function $P : \mathbb{R}^+ \rightarrow [0, 1]$ such that $P(r_i)$ is the probability that i successfully completes the task if it devotes r_i resources to the problem. We assume that $P(\cdot)$ is common to all providers, that if a provider devotes *no* resources to the task then it will fail (that is, $P(0) = 0$), and that the more resources are devoted to the task, the more likely it is to successfully complete the task. Thus, $P(\cdot)$ is continuous, increasing and strictly concave with $P(r_i) \rightarrow 1$ as $r_i \rightarrow \infty$. Finally, we assume that the probability of success of any service provider depends only on its own resource allocation, and not on the success or failure of any other provider.

We model the costly resources of provider i with a cost function, $c_i : \mathbb{R}^+ \rightarrow \mathbb{R}$. We assume that $c_i(\cdot)$ is continuous, increasing and convex, and that $c_i(0) = 0$. In addition, we assume that for any two service providers i and j , if they have different cost functions, $c_i(r)$ and $c_j(r)$, these functions are non-crossing for $r > 0$, i.e., either $c_i(r) = c_j(r), \forall r$ or $c_i(r) \neq c_j(r), r > 0$. An example class of cost functions which

¹ However, the *costs* of these resources (explained below) can become arbitrarily large.




Fig. 1. Figure (a) is an example quality of service function. Figure (b) shows example cost functions.

satisfy these properties are *linear* cost functions, where $c_i(r_i) = K_i r_i$ for some constant $K_i \geq 0$. Figure 1 illustrates the structure of possible cost and quality of service functions.

Since the contracting agent cannot directly observe the amount of invested resources (but only whether a task failed or succeeded), it must create incentives for service providers to invest a certain amount of resources. To do so, the contracting agent uses a payment scheme whereby the payment depends on whether or not the task was successfully completed. In the case that service provider i is contracted, it is automatically paid β_i , and then, if the task is successfully completed, it receives an additional bonus α_i . We assume that α_i is always non-negative, but place no restrictions on β_i .

3 The Full Information Setting

In this section we study the problem of procuring service providers to complete the agent's task, when the agent has full information about the providers' cost functions. We first analyse the case where the agent contacts a single provider, before extending these results to the case where the agent procures services from multiple providers.

3.1 Single Service Provider Case

Since the providers are self-interested and autonomous, the agent is unable to force them to execute the task. Instead, the agent must provide appropriate incentives so that the providers will take on the task for the agent, and will invest appropriate levels of resources. Assume that the agent has selected provider i to execute the task. By setting the parameter α_i appropriately, the agent is able to induce any desired level of effort, \bar{r}_i , from provider i .² In particular, for any value of \bar{r}_i , provider i 's expected utility is given by $U_i(\bar{r}_i) = \alpha_i P(\bar{r}_i) + \beta_i - c_i(\bar{r}_i)$ which is maximized when $U_i'(\bar{r}_i) = \alpha_i P'(\bar{r}_i) - c_i'(\bar{r}_i) = 0$. Solving for α_i , we have:

$$\alpha_i = \frac{c_i'(\bar{r}_i)}{P'(\bar{r}_i)} \quad (1)$$

² From here on, we use \bar{r}_i to denote the *agent's* desired level of effort, to distinguish from the *provider's* chosen investment r_i , and to emphasize that the agent cannot directly enforce this.

which is well defined (since $c_i(\cdot)$ is convex and $P(\cdot)$ is strictly concave), and is positive (since both $c_i(\cdot)$ and $P(\cdot)$ are increasing). Therefore, if the agent desires that i invest \bar{r} resources, then by setting $\alpha_i = c'_i(\bar{r})/P'(\bar{r})$, i maximizes its utility by actually investing \bar{r} (i.e. $\bar{r} = \operatorname{argmax}_r U_i(r)$).

As well as inducing effort, the agent would like to minimise the payment to provider i . At the same time, however, the agent needs to ensure that i will voluntarily participate, and as such must ensure that $U_i(\bar{r}_i) \geq 0$. The agent can satisfy both conditions by ensuring that $\beta_i \geq c_i(\bar{r}_i) - \alpha_i P(\bar{r}_i)$. By substituting α_i from Equation 1 and minimising β_i subject to the constraint, we have:

$$\beta_i = c_i(\bar{r}_i) - \frac{c'_i(\bar{r}_i)}{P'_i(\bar{r}_i)} P(\bar{r}_i). \quad (2)$$

While the agent wants to ensure that the selected provider invests the desired level of resources into the task, its real goal is to maximise its own utility, $U_A(\bar{r}) = (V - \alpha_i)P(\bar{r}) - \beta_i$. To this end, the agent must find α_i^* and β_i^* so as to induce the optimal level of effort from i , \bar{r}_i^* , such that $U_A(\bar{r}_i^*)$ is maximised. By substituting Equations 1 and 2 into the expression for the agent's utility, we get $U_A(\bar{r}_i) = VP(\bar{r}_i) - c_i(\bar{r}_i)$. Hence, by taking the first derivative and setting it to zero, and by letting \bar{r}_i^* denote the optimal level of (induced) investment by provider i from the agent's perspective, we get $V = c'_i(\bar{r}_i^*)/P'(\bar{r}_i^*)$ and therefore:

$$\alpha_i^* = V = \frac{c'_i(\bar{r}_i^*)}{P'(\bar{r}_i^*)} \quad (3)$$

To calculate β_i^* , let $g(r_i) = c'_i(r_i)/P'(r_i)$. Then $\bar{r}_i^* = g^{-1}(V)$ and thus:

$$\beta_i^* = c_i(g^{-1}(V)) - VP(g^{-1}(V)). \quad (4)$$

It is important to note that when V (and thus α_i) is very small, and the provider's marginal costs $c'(r)$ are very high, the optimal level of effort $\bar{r}_i^* = g^{-1}(V)$ may be negative. If this is the case, i cannot hope to obtain any (strictly) positive utility, irrespective of the actual invested effort. However, due to voluntary participation, the provider will then choose to not execute the task, and this is actually beneficial to the agent since costly providers will self-select and voluntarily opt out.

3.2 Multiple Service Providers Case

By procuring from multiple service providers, the agent may be able to increase the probability with which the task is successfully executed since if one of them fails, then another may succeed. We now consider two procurement strategies involving multiple service providers: (i) simultaneous procurement and (ii) sequential procurement (where the agent awaits the failure of one provider before approaching another).

Simultaneous Procurement. In simultaneous procurement, the agent contracts some subset of providers, M with $|M| = m \leq n$, to execute the task T . Once contracted, the m providers all execute the task at the same time, not waiting to see how others

perform. It may so happen that while the agent only requires the task to be completed once, multiple providers successfully complete the task. Given this possibility, the agent is faced with the problem of determining how to set the parameters α_i and β_i for each of the m providers in order to maximise its own utility.

Let $\bar{\mathbf{r}} = (\bar{r}_1, \dots, \bar{r}_n)$ be the resource vector specifying (induced) resource allocations for each provider i . Assume that $r_i = 0$ for all $i \notin M$. The expected utility of the agent, given that it selected the m providers to procure from, is:

$$\begin{aligned} U_A^{\text{sim}}(\bar{\mathbf{r}}, M) &= V \left(1 - \prod_{i \in M} [1 - P(\bar{r}_i)] \right) - \sum_{i \in M} [\alpha_i P(\bar{r}_i) + \beta_i] \\ &= V \left(1 - \prod_{i \in M} [1 - P(\bar{r}_i)] \right) - \sum_{i \in M} c_i(\bar{r}_i). \end{aligned} \quad (5)$$

Here, the latter equation is obtained by replacing α_i and β_i with Equations 1 and 2. By doing so, the agent induces investment \bar{r}_i , and in addition ensures voluntary participation. To determine the optimal investment levels for each provider i , \bar{r}_i^* , from the agent's perspective, we must solve $\nabla U_A^{\text{sim}}(\mathbf{r}, M) = 0$. This results in:

$$\alpha_i^* = V \prod_{j \neq i \in M} (1 - P(\bar{r}_j)). \quad (6)$$

The optimal solution $\bar{\mathbf{r}}^*$ and α_i^* for all i is found by solving the system of equations characterised by Equation 6, subject to the constraint that $\bar{r}_i^* \geq 0$ and $U_i(\bar{r}_i^*) = \alpha_i P(\bar{r}_i^*) + \beta_i^* - c_i(\bar{r}_i^*) \geq 0$. We note that, although at first glance it appears as though α_i^* does not depend on $c_i(\cdot)$, due to the interaction between the constraints for all providers, α_i^* is indirectly dependent on $c_i(\cdot)$. While this observation is interesting but irrelevant in the case where the cost functions are known, it will become important when we consider a mechanism for the incomplete information case in the next section.

Sequential Procurement. In sequential procurement, the agent only contracts an additional provider if previously contracted providers have failed to complete the task. This approach can be advantageous to the agent as it may only need to contract a small number of service providers, compared to the simultaneous procurement setting. On the other hand, the agent must decide the order in which to contract with providers.

The problem of sequentially procuring tasks from a set of service providers can be formulated as an *optimal search problem* [9]. If a provider i is contracted by the center and devotes r_i resources to the task, then it may successfully complete the task with probability $P(r_i)$, providing the agent with value $V - \alpha_i - \beta_i$, or the provider may fail with probability $1 - P(r_i)$, providing the agent with value $-\beta_i$. Weitzman showed that in such a setting, each provider can be characterized by an index given by:

$$z_i = \frac{P(r_i)(V - \alpha_i) - \beta_i}{P(r_i)} \quad (7)$$

The optimal order in which the providers should be procured corresponds to the decreasing order of z_i . That is, the provider with the largest index should be contracted

first, followed by the provider with the second largest index, if the first one fails, and so on. Given equations 1 and 2 which specify α_i and β_i , we have $z_i = V - c_i(r_i)/P(r_i)$. Since the cost functions of the providers are non-crossing, and quality of service function, $P(\cdot)$, is the same for all providers, then for a given r , providers with lower cost functions should be procured before those with higher ones.

Theorem 1. *Let M be a set of service providers, $|M| = m$, and assume $c_i(r) \leq c_{i+1}(r)$ for all r . Then when procuring the providers sequentially, i should be contracted before $i + 1$.*

Proof. Proof follows from [9].

Given that we now know the order in which to contract the providers, we can derive the utility of the agent. Assume that there are m providers, and that they are ordered such that $c_i(r) \leq c_{i+1}(r)$ for all r . The expected utility of the agent, given resource-allocation vector $\bar{\mathbf{r}} = (\bar{r}_1, \dots, \bar{r}_m)$ is:

$$U_A^{\text{seq}}(\bar{\mathbf{r}}) = U_A(\bar{r}_1) + \sum_{i=2}^m U_A(\bar{r}_i) \prod_{j=1}^{i-1} (1 - P(\bar{r}_j)) \quad (8)$$

In order to maximize the utility of the agent, set $\nabla U_A^{\text{seq}}(\bar{\mathbf{r}}) = 0$. This results in $m^* = V$ and:

$$\alpha_i^* = V - \left[U_A(\bar{r}_{i+1}) + \sum_{j=i+2}^m U_A(\bar{r}_j) \prod_{k=i+1}^{j-1} (1 - P(\bar{r}_k)) \right] \quad (9)$$

for $1 \leq i < m$. Unlike the simultaneous procurement case, α_i^* depends only on the cost functions of providers that are potentially contracted if i fails to complete the task. As was done in the single provider setting, to calculate β_i^* , we define $g_i(r) = c'_i(r)/P'(r)$, and then the optimal resource demand is $\bar{r}_i^* = g_i^{-1}(\alpha_i^*)$. Thus $\beta_i^* = c_i(\bar{r}_i^*) - \alpha_i^* P(\bar{r}_i^*)$.

4 Eliciting Cost Functions

In the previous section we assumed that the cost functions of the providers were known to the agent. Given this information, the agent was able to compute the appropriate values for parameters α_i and β_i so as to incentivise the providers in such a way that the agent's own expected utility was maximised. However, the assumption that the agent has full information is unrealistic in many practical applications. In this section, therefore, we show that it is possible to relax this assumption. In particular, we introduce a family of *task procurement mechanisms* $\text{TPM}(\alpha, \beta)$, such that providers are willing to truthfully reveal their cost functions to the agent, that then uses this information in order to set α_i and β_i appropriately.

Specifically, the agent first decides on a maximum number of providers, m ($1 \leq m \leq n$), it wants to procure services from (the optimal value of m can be determined experimentally, as shown in Section 5). The agent then executes $\text{TPM}(\alpha, \beta)$, which proceeds as follows:

1. *Cost elicitation:* All providers $i \in \{1, \dots, n\}$ report their cost functions $\hat{c}_i(\cdot)$ to the agent. We do not assume that providers reveal their true cost functions.

2. *Service provider selection and payment specification:* The agent selects the m providers with the lowest reported cost functions. We denote the set of chosen providers by M . Since we assume that the cost functions are non-crossing, there is no ambiguity in this selection. The agent then calculates α_i and β_i for each provider $i \in M$, and reports these parameters to the providers.
3. *Task execution:* If the agent uses a simultaneous procurement strategy, then all providers in M are asked to perform the task. If the agent uses a sequential procurement strategy then one provider, $i \in M$, is chosen at random to perform the task (Section 4.2 explains why this randomisation is important). If provider i fails, then another provider $j \in M \setminus \{i\}$ is chosen at random. This process continues until either a provider successfully completes the task, or all providers in M have attempted the task once and failed. Note that a provider is always allowed to refuse to attempt the task.
4. *Payment:* Any provider $i \in M$ that was contracted by the agent and successfully completed the task is paid $\alpha_i + \beta_i$. If provider i failed at the task then it receives β_i . All providers in M that were not asked to attempt the task, and those not in M initially, receive zero.

In the rest of this section we describe how to calculate parameters α_i and β_i so that the mechanism is incentive compatible and individually rational.

4.1 Simultaneous Procurement

We start by noting that the two parameters α and β allow us to define a *family* of incentive compatible mechanisms. Then, Theorem 2 characterises the family of such mechanisms for the simultaneous procurement strategy. We let \hat{c}_{m+1} denote the $(m+1)^{th}$ lowest reported cost function, and $r^*(\alpha, c)$ is the optimal investment decision of a provider with a cost function $c(\cdot)$ and when the agent announces parameter α .

Theorem 2. *If, for all $i \in M$, α_i is independent of $\hat{c}_i(\cdot)$, and β_i is given by:*

$$\beta_i = \hat{c}_{m+1}(r^*(\alpha_i, \hat{c}_{m+1})) - VP(r^*(\alpha_i, \hat{c}_{m+1})) \quad (10)$$

then $TPM(\alpha, \beta)$ with simultaneous procurement is individually rational and incentive compatible.

Proof. Individual rationality holds because a provider i can always refuse the task after learning α_i and β_i . Since all providers in M are asked to attempt the task, the expected utility of $i \in M$ does not depend on the payment or the resources allocated to the task by any other providers $j \in M$. Thus we can look at the incentives of each provider $i \in M$ independently. Consider first provider i with cost function c_i . If $c_i < \hat{c}_{m+1}$ then i has no incentive to announce $\hat{c}_i > \hat{c}_{m+1}$ since by doing so it would guarantee itself a utility of zero. If it revealed its cost function truthfully, then its expected utility is greater than or equal to zero. Now consider the case where provider i has true cost function $c_i > \hat{c}_{m+1}$, but misreports $\hat{c}_i < \hat{c}_{m+1}$ and is selected. Let $U_i(r, c) = \alpha_i P(r) + \beta_i - c(r)$. Note that β_i is set such that $U_i(r^*(\alpha_i, \hat{c}_{m+1}), \hat{c}_{m+1}) = 0$. It follows that $U_i(r^*(\alpha_i, c_i), \hat{c}_{m+1}) < 0$, and since $c_i > \hat{c}_{m+1}$, $U_i(r^*(\alpha_i, c_i), c_i) < U_i(r^*(\alpha_i, c_i), \hat{c}_{m+1}) < 0$. Thus such a

provider always receives negative utility. Together with the fact that payments to the providers are independent of their reports, i has no incentive to under-report its cost function.

We now introduce two mechanisms that satisfy the above requirement for incentive compatibility: (1) *uniform pricing* and (2) *discriminatory pricing*. From Theorem 2, since β is given by Equation 10, we only need to worry about setting α_i for all $i \in M$. Now, in the uniform pricing mechanism, $\alpha = \alpha_1 = \dots = \alpha_m$, and we calculate α by replacing \bar{r}_j in Equation 6 by $r^*(\alpha, \hat{c}_{m+1})$, resulting in following equation:

$$\alpha = V [1 - P(r^*(\alpha, \hat{c}_{m+1}))]^{m-1} \quad (11)$$

Corollary 1. *Uniform Pricing Mechanism. TPM(α, β) with α and β satisfying Equations 11 and 10 is individually rational and incentive compatible.*

Proof. Since α is independent of any of the reports of providers in M , the proof follows from Theorem 2.

Although uniform pricing is a natural extension of the single-provider case, there are better alternatives. In particular, although from Theorem 2, α_i needs to be independent of \hat{c}_i , we can use the cost functions of other providers to calculate α_i . In fact, from Equation 6 the optimal α_i^* in the complete information case can be calculated from $\bar{r}_j, j \neq i$ and thus does not directly depend on \hat{c}_i . However, as noted in Section 3.2, if we solve optimally for i as well as all j simultaneously, then α_i^* depends indirectly on \hat{c}_i since \bar{r}_j^* depends on \bar{r}_i , and r_i^* in turn depends on \bar{r}_j^* . Nevertheless, we can take advantage of the information available about the cost functions of other providers to develop a discriminatory pricing mechanism, though we need to be careful about how we calculate α_i to ensure incentive compatibility. In more detail, we calculate α_i for a specific provider $i \in M$ by solving the following system of equations:

$$\alpha_i = V \prod_{j \in M \setminus \{i\}} [1 - P(r^*(\alpha_j, \hat{c}_j))] \quad (12)$$

and for all $j \in M \setminus \{i\}$:

$$\alpha_j = V [1 - P(r^*(\alpha_i, \hat{c}_{m+1}))] \prod_{k \in M \setminus \{i, j\}} [1 - P(r^*(\alpha_k, \hat{c}_k))] \quad (13)$$

Note that we need to derive and solve a separate set of equations for each $i \in M$.

Corollary 2. *Discriminatory Pricing Mechanism. If, for each $i \in M$, α_i is given by independently solving equations 12 and 13, and if β_i is given by Equation 10, TPM(α, β) is incentive compatible and individually rational.*

Proof. Note that \hat{c}_i does not appear in equations 12 and 13, and thus α_i is independent of that provider's report. The proof follows directly from Theorem 2.

4.2 Sequential Procurement

In the full-information setting, described earlier, the order in which the providers are asked to perform the task was shown to be important. However, in the incomplete information setting, we cannot use the reported costs to determine this order since it influences the expected utility of the providers (since those with higher cost functions are less likely to be contracted by the agent). For this reason, the providers are randomly selected in the third stage of the mechanism, and we refer to $i \in M$ as the i^{th} provider in the random sequence, but \hat{c}_{m+1} is the $(m+1)^{\text{th}}$ lowest reported cost as before.

Now, whereas in the previous section we were able to use the reported costs of the other providers to adjust the payment without changing the incentives, this is no longer the case in the sequential setting. To see this, note that from Equation 9, α_i^* depends on all providers that appear after i in the sequence. Given these considerations, Theorem 3 reformulates the requirements for incentive compatibility in terms of the sequential procurement setting.

Theorem 3. *If, for all $i \in M$, α_i is independent of any $\hat{c}_j(\cdot)$, $m \geq j \geq i$, and β_i is given by Equation 10, then $\text{TPM}(\alpha, \beta)$ with sequential procurement is individually rational and incentive compatible.*

Proof. Since $\hat{c}_i(\cdot)$ does not affect the payment of providers $j > i$, the quality of service offered by these providers is independent of this report. Furthermore, the position of provider i in the sequence is independent of $\hat{c}_i(\cdot)$. As a result, other than through the allocation decision, provider i 's utility is independent of $\hat{c}_i(\cdot)$. Furthermore, analogous to the arguments in Theorem 2, β_i is set such that providers cannot benefit by misreporting in order to change the allocation.

We now present a specific discriminatory payment scheme where $\alpha_m = V$, and α_i for $i < m$ is calculated by modifying Equation 9 to give:

$$\alpha_i = V - \left[U_A(r^*(\alpha_{i+1}, \hat{c}_{m+1})) + \sum_{j=i+2}^m U_A(r^*(\alpha_j, \hat{c}_{m+1})) \prod_{k=i+1}^{j-1} [1 - P(r^*(\alpha_k, \hat{c}_{m+1}))] \right] \quad (14)$$

Corollary 3. *Sequential Procurement. $\text{TPM}(\alpha, \beta)$ where $\alpha_m = V$, α_i for $i < m$ is given by Equation 14, and β_i is given by Equation 10, is individually rational and incentive compatible.*

Proof. Since α_i is independent of any reports of providers in M , the proof follows directly from Theorem 3.

5 Empirical Evaluation

Having described three incentive compatible and individually rational mechanisms, we now instantiate the quality of service and cost functions, and empirically evaluate our approach. The purpose of this evaluation is two-fold. First, it provides an example of a task allocation domain and how the mechanisms can be applied. Second, it allows us to




Fig. 2. Simulation results showing the expected utility of the agent for simultaneous procurement with uniform pricing (squares), simultaneous procurement with discriminatory pricing (triangles) and sequential procurement (circles). Results for optimal sequential (short dashed) and simultaneous (long dash) procurement with complete information are also shown.

compare and evaluate the different mechanisms which we presented against the original case in which a single U_A provider is contracted³. In more detail, we define the quality of service function as $P(r) = 1/(1 + 1/r)$ and the providers' costs as linear functions $c_i(r) = K_i r$. As required, $P(\cdot)$ is strictly concave, $P(0) = 0$ and $P(r) \rightarrow 1$ as $r \rightarrow \infty$. Furthermore, note that $c(\cdot)$ is convex (although not strictly convex) and $c(0) = 0$. In the simulations that follow the constants, K_i , are independently drawn from a uniform distribution with support $[1, 2]$.

We choose these functions since they are representative of the general class of functions to which our formalism applies, and also because they yield attractive analytical solutions. For example, each provider maximises its utility by committing resources $r_i^* = \sqrt{\alpha_i/K_i} - 1$ and this will result in a quality of service such that $P(r_i^*) = 1 - \sqrt{K_i/\alpha_i}$. Furthermore, Equation 6 becomes $\alpha_i^* = V \prod_{j \neq i \in M} \sqrt{K_j/\alpha_j^*}$ and can be solved by taking the logarithm of both sides, and using the standard numerical technique of Gauss-Seidel iteration [10], incorporating the voluntary participation constraint.

Figure 2 shows the results of applying the simultaneous (with uniform and discriminatory pricing) and sequential procurement mechanisms in this setting. We vary both the total number of providers, n , and the size of the subset of them that are selected by the agent, m , and in all cases $V = 4$. We first note that the expected utility of the agent when using the simultaneous procurement mechanism with discriminatory pricing always exceeds that of the mechanism with uniform pricing, and thus discriminatory pricing is always preferred. This is not surprising, since discriminatory pricing makes more use of the cost information that is available to the agent. Likewise, the expected utility of the agent when using the sequential procurement mechanism always exceeds that of either simultaneous procurement mechanism, and thus, if sequential procurement is feasible in the specific application domain, then it is always preferred in both the full and incomplete information settings. Interestingly, the difference in utility between the full and incomplete information setting is much larger in the case of

³ Note that our simultaneous and sequential procurement mechanisms are identical in the case that $m = 1$, and since the single provider mechanism is efficient, this is also identical to Matsubara's.

the sequential procurement mechanism compared to either simultaneous procurement mechanism. Finally, the results show that given any specific setting (i.e. the valuation of the agent, the total number of providers and the distribution that describes their cost functions), there is an optimum number of providers to select to procure the task from (either simultaneously or sequentially), and we note that when n is large, procurement from multiple service providers is preferable to procuring from a single one, over a wide range of values of m .

6 Conclusions

In this paper, we considered the problem of procuring computational tasks from self-interested service providers that are able to flexibly manipulate their quality of service in order to maximise their own utility, and we derived a family of task procurement mechanisms that allowed a contracting agent to procure tasks from multiple service providers (either simultaneously or sequentially). Our future work in this area concerns extending these results to the case in which the agent wishes to procure multiple interdependent tasks which exhibit complementary and substitutable valuations. This setting corresponds to the problem of procuring services within a computational workflow, where the entire workflow may be worthless if particular tasks are not completed successfully. Such an extension is likely to require the use of combinatorial auctions, and we are particularly interested in exploring the existence of individually rational and incentive compatible mechanisms that are also optimal.

References

1. Foster, I., Jennings, N.R., Kesselman, C.: Brain Meets Brawn: Why Grid and Agents Need Each Other. In: Proc. of the 3rd Int. Conference on Autonomous Agents and Multiagent Systems, vol. 1, pp. 8–15 (2004)
2. Dash, R.K., Ramchurn, S.D., Jennings, N.R.: Trust-based mechanism design. In: Proc. of the 3rd Int. Conference on Autonomous Agents and Multi-Agent Systems, pp. 748–755 (2004)
3. Porter, R., Ronen, A., Shoham, Y., Tennenholtz, M.: Fault tolerant mechanism design. *Artificial Intelligence* 172(15), 1783–1799 (2008)
4. Matsubara, S.: Trade of a problem-solving task. In: AAMAS 2003: Proceedings of the second international joint conference on Autonomous agents and multiagent systems, pp. 257–264. ACM, New York (2003)
5. Miller, N., Resnick, P., Zeckhauser, R.: Eliciting informative feedback: The peer-prediction method. *Management Science* 51(9), 1359 (2005)
6. Papakonstantinou, A., Rogers, A., Gerding, E.H., Jennings, N.R.: A truthful two-stage mechanism for eliciting probabilistic estimates with unknown costs. In: Proc. of the 18th European Conference on Artificial Intelligence, Patras, Greece, pp. 448–452 (2008)
7. Stein, S., Jennings, N.R., Payne, T.: Flexible service provisioning with advance agreements. In: Proc. of the 7th Int. Conference on Autonomous Agents and Multi-Agent Systems, pp. 249–256 (2008)
8. Jehiel, P., Moldovanu, B.: Efficient Design with Interdependent Valuations. *Econometrica* 69(5), 1237–1259 (2001)
9. Weitzman, M.L.: Optimal search for the best alternative. *Econometrica* 47(3), 641–654 (1979)
10. Press, W.H., Flannery, B., Teukolsky, S.A., Vetterling, W.T., et al.: Numerical recipes. Cambridge University Press, New York (1986)