Flexible Contracting with Heterogeneous Agents and Stochastic Demand

D. Singham[1], W. Cai[2], and A. Fügenschuh[3]

[1]Operations Research Department
Naval Postgraduate School
Monterey, CA 93943
dsingham@nps.edu

[2]Department of Mechanical and Industrial Engineering
New Jersey Institute of Technology
Newark, NJ 07102
cai@njit.edu

[3]Institute of Mathematics
Brandenburg University of Technology
Platz der Deutschen Einheit 1, 03046 Cottbus, Germany
fuegenschuh@b-tu.de

August 29, 2018

Abstract

While the predominant pricing scheme for services is a set of long-term contracts stipulating a fixed fee for a level of usage, many electronic service providers now offer flexible contracts. These contracts include flexibility in usage level and the option to opt out of the service in certain time periods. We study the effectiveness of flexible contracts in aligning the interests between the service provider and the customers in the presence of customers with different discrete demand distributions. This work develops a new framework using the principal-agent framework, and is particularly relevant to e-commerce for service providers who could implement complex pricing mechanisms using these models. We explore two types of contract variations. The first is aggregated versus differentiated contracts, the second is whether agents may decline participation in a future time period after committing to a set of contract options. We find that the principal always prefers differentiation to aggregation, and under differentiation prefers to require participation at each time period. Alternatively, the agents' preference depends on their value function and demand distribution. Our study provides insights that help decision makers price their services to better satisfy customers' varying needs while improving the profitability of the service.

Keywords: principal-agent model; contracts; heterogeneous agents; stochastic demand.

1 INTRODUCTION

The principal-agent (PA) framework can be used to determine the quantity and price for a service that a principal should offer an agent when the agent has hidden information regarding his performance, costs, or preferences. This paper considers the setting where a principal maximizes her expected profit while facing an adverse selection problem regarding the agent's demand for a product or a service. In particular, the agents have heterogeneous demand distributions, and the principal does not know the agent's demand distribution prior to contracting. After contracting, both the principal and the agent face demand uncertainty at each time period according to the agent's demand distribution. We evaluate different contract frameworks to handle these types of agents.

The proposed model differs from the classical PA framework by considering information asymmetry both *ex ante* and *ex post*. Prior to contracting, while the agent knows his actual demand distribution, the principal only knows the possible demand distributions and the probability of an agent's demand following each distribution. The principal offers a menu of contracts, each contract designed for a particular demand distribution. The agent reveals his distribution through the selection of a contract. After contracting, the agent faces demand uncertainty associated with the agent's demand distribution in any given time period. Under the flexible contracting options we will propose, at each time period the agent can select an option from the pre-agreed contract. This differs from the classical PA model where the demand distribution is common knowledge to both parties and the particular demand value is assumed to be private knowledge which is realized before contracting.

The framework we present is a generalization of many PA settings in which information asymmetry exists in the contracting stage only. Here, we consider the effects of stochastic information post-contracting, which allows for potentially more flexible contracting. One example is contracting for information goods, which typically entail a fixed usage amount for a fee each month, and an agent picks his plan according to his usage. However, actual usage of information goods can vary month to month (for example, cell-phone data usage). Another example includes energy users who may face a stochastic demand for power each month after agreeing to a fixed payment contract. We will demonstrate when the principal and agent can benefit from flexible contracts that accommodate demand uncertainty post-contracting.

In addition to the relaxation of the assumptions on the information structure, this paper considers two possible variations on contract types that the principal can offer the agents with D possible heterogeneous discrete demand distributions, each with N possible demand levels. The first variation is that the principal can offer "*differentiated*" or "*aggregated*" contracts towards each of the D demand distributions. Differentiated contracts allow the agent to choose from N different options at each time period, so each of the D contracts has N individual options tailored to different demand realizations. Aggregated contracts only offer one option post-contracting which targets the entire demand distribution, analogous to a supply chain coordination contract in a newsvendor setting. Many cell-phone plans are aggregated contracts, in that once the agent chooses his plan according to his anticipated demand distribution, he obtains a fixed data usage limit each month for a fixed price. We find that the principal always prefers to offer differentiated contracts to aggregated contracts.

The second variation in contract type is "*forced participation*" versus "*nonparticipation options*." Forced participation means once an agent selects a contract, he will have to participate by choosing a non-zero quantity/price option at each future time period. Nonparticipation options allow the agent to decline participation in any future time period if none of the available options gives him a nonnegative utility. This is equivalent to a shutdown option always being available post-contracting. For example, some cell-phone plans allow minimal users to opt-in or out of a data plan each month depending on whether they anticipate needing data. We find that while agents generally prefer contracts with nonparticipation options because of their flexibility, there are some demand distributions where forced participation contracts may give the agent a higher expected utility, as the principal may need to lower prices to induce agents to enter into the contract in the first place. When participation is forced, principals prefer differentiated contracts to aggregated contracts.

We highlight the contributions of the paper relative to past work. Cai and Singham (2018) relax the traditional assumptions on the information structure in two ways: the agent knows his true demand distribution at contracting but not particular demand realizations at future time periods, and the principal does not know the probability distribution of the agent but can estimate a finite number of possible distributions. The work in the present paper makes similar relaxations in the information structure to Cai and Singham (2018), but has three new contributions to advance the literature. The first is that analysis and results are generalized in the present work while Cai and Singham (2018) consider a specific formulation relating to carbon capture and storage systems. The second contribution is that the formulations are designed for arbitrary dimensions (D and N are arbitrary) and employs more sophisticated numerical methods for solving these larger optimization problems. Cai and Singham (2018), on the other hand, present analytical solutions for the small-dimensional specific formulation presented, and only consider differentiated contracts with nonparticipation options. The third, and most important, contribution is that this paper considers the joint effects of differentiation/aggregation and forced/nonparticipation, allowing for improved flexibility and optimization over contracting choices uniquely designed to deal with heterogeneous agent information structures. In particular, we show when the principal and agent are likely to prefer certain contract types.

1.1 Pricing Services

One of the applications for the proposed framework is service pricing. The predominant scheme for services, such as cable, cell phone data usage, and e-content, is a schedule of fixed fees. A service provider offers a menu of pricing plans where each plan consists of a fixed fee for a level of usage or a combination of features. Customers choose plans that best suit their needs. Once the plans are selected, customers cannot change the level of usage or the number of features from one period to the next, and cannot skip certain periods. This type of contract is an example of *aggregated contracts with forced participation.*

Though charging customers a fixed fee provides (near) certainty to customers on their costs and predictability to service providers on their revenue, the heterogeneity among customers' consumption levels can lead to one of the following consequences. First, high-volume customers drive up the overall usage and the service providers must increase the fee in order to sustain the service quality. Second, low-volume users are left with two options: not using the service, or overpaying.

A usage-based pricing contract, unlike the fixed-fee pricing contract, charges customers based on the quantity consumed. This pricing scheme benefits those customers who consume small quantities and are unwilling (or unable) to pay fixed-fees. Realizing the benefits, some firms have combined the fixed-fee pricing with a usage-based, linear pricing scheme. MetroMile car insurance charges customers “a low base rate plus a few cents per mile” up to a certain mileage per day. Similarly, Amazon Chime provides voice and video calls with no subscription fees. A $3 fee is charged to the meeting host for each usage day and the total fee is capped at a maximum of $15 per month. Such practice aligns well with the theoretical results presented in Sundararajan (2004), who shows that the optimal pricing menu has the following three characteristics: an unlimited usage fixed-fee contract will be offered towards the customers who have the highest demand types, a usage-based contract will be offered toward those whose types are in the mid-range, and the customers with the lowest demand types will not be served.

Technological advancement enables the execution of usage-based, nonlinear pricing schemes which have become increasingly popular in digital services. Jelastic, a company that provides cloud services for hosting providers or application developers, adopts the usage-based pricing model. A hosting provider or an application developer selects a minimum resource level, i.e., the number of cloudlets, which then determines the base cost. Jelastic gives a nonlinear volume discount based on the minimum resource level. As more customers use the application, more cloudlets are added automatically to meet the demand. The hosting provider is then charged based on usage and the minimum resource level. This is an example of *differentiated contracts with forced participation.*

There are also multiple examples of flexible contracts with non-participation options. One example is Freshly, a meal delivery service. Users can opt in or out each week and can also vary their quantity (number of meals delivered) each week. The pricing scheme is nonlinear in quantity. This type of pricing is an example of *differentiated contracts with non-participation options.* Another example includes cell-phone plans that allow users to opt out of data entirely in the base plan, but opt in to a capped data amount for a fee on a week by week basis. In this case, the contract is an example of *aggregated contracts with non-participation options.*

The following describes the layout of the paper. Section 2 presents the relevant literature. Section 3 describes four models: differentiated contracts with forced participation (3.1.1) and nonparticipation (3.1.2), and aggregated contracts with forced participation (3.2.1) and nonparticipation (3.2.2). Section 4 presents managerial results. Section 5 delivers numerical results from optimal contracts over a variety of different value functions and demand distributions. Finally, Section 6 concludes.

2 LITERATURE REVIEW

We present two pertinent areas of literature in this section. The first stream of literature is the adoption of the principal-agent framework in supply chain, service, and digital goods contracting. Though this framework can deal with two types of information asymmetry: adverse selection and moral hazard, we focus our review on the former. The second body of literature is on computational methods for solving complex decision models.

2.1 Applications of the Principal-Agent Framework

In the classical principal-agent framework, agents are assumed to privately observe their types in either valuation, quality or demand. Facing such information asymmetry, the principal designs incentives to induce desired behavior from agents anticipating their preferences. Maskin and Riley (1984) show that a nonlinear price-quantity schedule can discriminate among a set of buyers with discrete types under information asymmetry as long as the single-crossing property holds.

This framework has been widely adopted in the literature that evaluates the effectiveness of various contracts that coordinate decentralized supply chains. Corbett et al. (2004) find that a per unit price and a side payment as a function of ordering quantity outperforms linear contracts in obtaining the buyer's cost information. Iyer et al. (2005) show that a buyer can screen the supplier's capability by designing a menu of contracts that stipulate the buyer's resource-commitment and the transfer price to the supplier. Özer and Wei (2006) demonstrate that a supplier can elicit the manufacturer to disclose his private forecast information by offering a menu of nonlinear capacity reservation contracts, which specify the payment and capacity as functions of forecasts. Similarly, it is shown by Çakanyildirim et al. (2012) that the supplier can offer a menu of contracts that specify the ordering quantity and the supplier's share of the channel profit to the retailer to coordinate the supply chain even though the retailer does not observe the supplier's production cost. Burnetas et al. (2007); Taylor and Xiao (2009); Yang et al. (2009); Babich et al. (2012); Yang et al. (2012); Kim and Netessine (2013); Chen et al. (2016); Tran and Desiraju (2017) address the incentive misalignment due to information asymmetry between key parties (such as buyer-supplier, manufacturer-retailer, and manufacturer-supplier) by designing contracts that improve channel coordination and its profitability using the PA framework.

Another application area of the principal-agent model is to tackle the adverse selection problem in service contracting. Akan et al. (2011) study the optimal contract for outsourcing service providers when the outsourcing firm has private information about the demand level. Jiang et al. (2012) argue that linear performance-based contracts perform just as well as the fee-for-service contract to reduce delays in patient services when hospital capacity allocations are observable and contractible. More closely related to our work, Sundararajan (2004) studies digital goods pricing when customers have heterogeneous consumption quantities. The optimal pricing schedule includes a fixed-fee contract for high-usage customers and a usage-based contract for medium-usage customers. Turning attention from quantity to quality, Bhargava and Choudhary (2008) show that offering different qualities of information goods at different prices, or versioning, is better than offering a single quality when the firm's demand from selling the low quality alone exceeds that from selling the high quality alone.

Many of the studies adhere to the information structure of the classical principal-agent problem. That is, the agent learns his type (demand, cost, capacity, etc.,) *before* accepting any incentives and there is no uncertainty in its value. In the case where the value changes from one period to the next, a dynamic adverse selection model is then developed. Arya and Mittendorf (2006) show that in a memory-based contract, the firm's funding decision not only depends on the manager's cost report in the second period, but also on the report in the first period, and can outperform contracts where funding decisions are based on only the current period's cost report. Zhang et al. (2010) argue that offering a stationary single quantity for a fixed price can outperform the base-stock policy for the supplier when the retailer's inventory at the beginning of each period is unobservable. Löffler et al. (2012) compares a buyer's supplier decision under hurried contracting with an abandonment option, where contracting occurs before the new supplier learns his production cost, and under delayed contracting, where contracting occurs after cost information is revealed. It is shown that the buyer is more likely to purchase from the new supplier under hurried contracting as she will have to compensate the new supplier with a payment contingent on the reported production cost if she chooses to use the existing supplier. Chen and Deng (2015) conclude that a supplier may voluntarily reveal his privately observed production efficiency to the manufacturer when such information is imprecise. This then leads to an increase in production quantity, and may even result in over-production.

2.2 Computational Methods

There has been an increasing volume of literature on solving decision models without the assumption of full information about uncertainty. Bertsimas and Thiele (2006) develop a framework that uses observations of random variables as direct inputs to mathematical programming problems, thus allowing the use of historical data without making any assumption on the distribution of uncertainty. Perakis and Roels (2010) show that in the content of network revenue management, this approach not only obtains an efficient solution but also can incorporate information about demand into an uncertainty set at no modeling or computational cost. Ehtamo et al. (2010) propose an online adjustment scheme when the buyer's value function (instead of just the type) is not known. The seller may improve his profit and the buyers' utilities in a dynamic game setting by adding a new price and quantity bundle in addition to the old one in each period.

One disadvantage of the principal-agent framework is the lack of tractable solutions when the agent's type follows a discrete distribution with more than two values. As Lambert (2006) points out, relying on closed-form results limit the type and complexity of models that can be solved. Following the proposed method by Dempe (1995), Cecchini et al. (2013) formulate principal-agent problems as bilevel nonlinear programs and solve them numerically. Singham (2018) finds approximations to continuous principle-agent problems with non-tractable formulations using discrete approximations with a large number of agent types.

3 MODEL FORMULATIONS

We start with a general framework for optimal contracts for the principal who wants agents to participate in selecting contracts. There are D possible different demand distributions for the agents. Each distribution is denoted θ_d and the agent has probability μ_d of having distribution $d, d = 1, \ldots, D$ with $\mu_d \geq 0$ and $\sum_{d=1}^{D} \mu_d = 1$. Let $\mathcal{D}$ be the set $\{1, \ldots, D\}$. Let x_d be a contract (vector of options) designed by the principal targeted towards an agent with distribution d, where x_d consists of quantities of a good to be transferred q_d, and a total price charged t_d. The expression $\Phi_d(x_d, \theta_d)$ represents the expected profit to the principal in a single time period if an agent with demand distribution θ_d selects contract x_d. Finally, let $u(x_{d'}, \theta_d)$ be the utility to an agent with demand distribution θ_d selecting contract $x_{d'}$. The optimal profit formulation under heterogeneous demand distributions is:

$$
\begin{array}{rlllll}
\Phi = & \max\limits_{x_d} & \sum_{d=1}^{D} \mu_d \Phi_d(x_d, \theta_d) & & & \\
& \text{s.t.} & u(x_d, \theta_d) \geq 0 & d \in \mathcal{D} & (IR_d) & \\
& & u(x_d, \theta_d) \geq u(x_{d'}, \theta_d) & d, d' \in \mathcal{D} & (IC_{dd'}) & \quad (1)\\
& & q_d \geq 0, t_d \geq 0 & d \in \mathcal{D} & (NN_d). &
\end{array}
$$

For brevity, the indices for the decision variables will be implied throughout the paper, i.e., instead of $x_d, d \in \mathcal{D}$ we just write x_d. The objective function is a weighted average of the expected profit for each possible demand distribution. The IR_d (*individual rationality*) constraints ensure that agents receive nonnegative utility from participating. The $IC_{dd'}$ (*incentive compatibility*) constraints ensure they will select the contract designed for their distribution. Constraint NN_d ensures that the quantity and prices offered in each option are nonnegative. In some special cases, t_{dn} can be negative if the principal offers a subsidy to the agent, but we constrain it to be nonnegative here for simplicity. For each of the following models, we introduce detailed notation and will define specific forms for $\Phi_d(x_d, \theta_d)$ and $u(x_{d'}, \theta_d)$.

3.1 Differentiated Models

The differentiated model occurs when the principal offers different contract options towards each possible demand level within each demand distribution. We next describe the optimal formulation for differentiated models with forced participation (after agents select a contract at the initial contracting stage, agents must select an option each time period), and with nonparticipation options (agents may choose not to participate in any period after contracting depending on their observed demand).

3.1.1 *Differentiated Contracts with Forced Participation*

There are two dimensions of uncertainty in demand θ, in that there are multiple possible discrete distributions instead of just one distribution. Without loss of generality, we assume each demand distribution θ_d contains N discrete nonnegative demand levels $(n = 1, \ldots, N)$ in increasing order. We note that dummy demands of zero can always be used to pad a distribution so it has N levels. Let $\mathcal{N}$ be the set $\{1, \ldots, N\}$. The left matrix of (2) shows the different possible values of θ_{dn}. The principal now designs a **menu** of **contracts** (rows), with each row containing **options** designed to address the uncertainty in that given demand distribution. These options are displayed in the right-hand side of (2), with t_{dn} as the total price paid to receive q_{dn} units.

$$\begin{bmatrix} \theta_1 \\ \vdots \\ \theta_d \\ \vdots \\ \theta_D \end{bmatrix} = \begin{bmatrix} \theta_{11} \cdots \theta_{1n} \cdots \theta_{1N} \\ \vdots \\ \theta_{d1} \cdots \theta_{dn} \cdots \theta_{dN} \\ \vdots \\ \theta_{D1} \cdots \theta_{Dn} \cdots \theta_{DN} \end{bmatrix} \Longrightarrow \begin{bmatrix} (q_{11}, t_{11}) \cdots (q_{1n}, t_{1n}) \cdots (q_{1N}, t_{1N}) \\ \vdots \\ (q_{d1}, t_{d1}) \cdots (q_{dn}, t_{dn}) \cdots (q_{dN}, t_{dN}) \\ \vdots \\ (q_{D1}, t_{D1}) \cdots (q_{Dn}, t_{Dn}) \cdots (q_{DN}, t_{DN}) \end{bmatrix} \tag{2}$$

In this setting, there are two types of information asymmetry. At the contracting stage, $t = 0$, the agent knows his type d and both parties know μ_d. The principal provides a menu of contracts, each contract designed towards a particular distribution d. The agent selects a contract (a row of the RHS of (2)) thus revealing his distribution d. For the remaining time periods $t = 1, 2, \ldots,$ the agent privately observes different demand realizations according to his distribution $\pi_d = [\pi_{d1}, \ldots, \pi_{dn}, \ldots \pi_{dN}]$, where $\pi_{dn} \geq 0$ and $\sum_{n=1}^{N} \pi_{dn} = 1$. The options within the contract $\begin{bmatrix} (q_{dn}, t_{dn})_{n=1,\ldots,N} \end{bmatrix}$ will induce the agent to reveal his true demand. Each time period he selects the option dn from the previously selected contract (row d) that maximizes his utility. The principal then delivers q_{dn} units and receives payment t_{dn} from the agent.

The principal's objective is to choose the contract options in the RHS of (2) at $t = 0$ so that her expected profit is maximized over times $t = 1, 2, \ldots,$ etc. Let $v(q_{d'n'}, \theta_{dn})$ be the value function of an agent obtaining $q_{d'n'}$ units of a product when his demand is θ_{dn}. His utility associated with choosing option $(q_{d'n'}, t_{d'n'})$ is $v(q_{d'n'}, \theta_{dn}) - t_{d'n'}$. Let $x_{d'}$ denote the row associated with contract d', so that $x_{d'} = [(q_{d'1}, t_{d'1}) \cdots (q_{d'n}, t_{d'n}) \cdots (q_{d'N}, t_{d'N})]$. If an agent with demand distribution θ_d chooses contract

$x_{d'}$ at $t = 0$, his expected utility given that he will choose the options that maximize his utility in future time periods is

$$u(x_{d'}, \theta_d) = \sum_{n=1}^{N} \pi_{dn} \max_{n'} \{v(q_{d'n'}, \theta_{dn}) - t_{d'n'}\}. \tag{3}$$

Now suppose the principal has a cost function $s(q_{dn}, \theta_{dn})$ associated with delivering quantity q_{dn} to an agent with demand θ_{dn}. If an agent with demand distribution θ_d correctly chooses contract x_d, and will choose the correct option corresponding to his demand in future time periods, the expected profit to the principal is

$$\Phi_d(x_d, \theta_d) = \sum_{n=1}^{N} \pi_{dn}(t_{dn} - s(q_{dn}, \theta_{dn})). \tag{4}$$

We will use the superscript DF to denote the differentiated model with forced participation, and will use Φ^{DF} to refer to both the formulation and the optimal objective value as needed. The principal's optimization problem Φ^{DF} is

$$\begin{array}{llllll}
\Phi^{DF} = & \max\limits_{\{q_{dn}, t_{dn}\}} & \sum_{d=1}^{D} \mu_d \sum_{n=1}^{N} \pi_{dn}(t_{dn} - s(q_{dn}, \theta_{dn})) & & & \\
 & \text{s.t.} & u(x_d, \theta_d) \geq 0 & d \in \mathcal{D} & (IR_d) & \\
 & & u(x_d, \theta_d) \geq u(x_{d'}, \theta_d) & d, d' \in \mathcal{D}, d \neq d' & (IC_{dd'}) & (5) \\
 & & v(q_{dn}, \theta_{dn}) - t_{dn} \geq v(q_{dn'}, \theta_{dn}) - t_{dn'} & d \in \mathcal{D}, n, n' \in \mathcal{N}, n \neq n' & (IC_{dnn'}) & \\
 & & q_{dn} \geq 0, t_{dn} \geq 0 & d \in \mathcal{D}, n \in \mathcal{N} & (NN_{dn}). &
\end{array}$$

The objective function is the expected profit to the principal across all distributions, assuming that agents with demand distribution d select contract x_d. The constraints ensure that agents select the contract designed for their distribution at $t = 0$, and that they will then select the option geared towards their realized demand level at each future time. The first two sets of constraints, taken from (1), are *individual rationality* and *incentive compatibility* constraints that ensure agents with distribution d select contract x_d at the contracting stage. We have an additional constraint $IC_{dnn'}$, which is the *incentive compatibility* constraint at the option level. This constraint ensures type-d agents with realized demand θ_{dn} choose option dn instead of the other options dn', $n' \neq n$. This constraint is needed to ensure that the π_{dn} probabilities are realized in the objective and utility functions. We now present the following assumptions for all models discussed in this paper.

Assumption 3.1. Assume that $v(q_{d'n'}, \theta_{dn})$ and $s(q_{d'n'}, \theta_{dn})$ are bounded, nonnegative, increasing functions of $q_{d'n'}$ and θ_{dn}. Additionally let $v(q_{d'n'}, \theta_{dn})$ be concave in $q_{d'n'}$, and $v(q_{d'n'}, \theta_{dn}) = 0$ and $s(q_{d'n'}, \theta_{dn}) = 0$ if $q_{d'n'} = 0$.

Assumption 3.2. Assume the single-crossing property holds, where the derivative of $v(q_{d'n'}, \theta_{dn})$ with respect to $q_{d'n'}$ is increasing in θ_{dn}.

Assumption 3.3. For any option (q, t), we have $t \leq qM$ for some large positive constant M. This means the principal cannot charge a positive price t while delivering $q = 0$.

These assumptions are standard in principal-agent theory, see Laffont and Martimort (2009) for an overview.

3.1.2 Differentiated Contracts with Nonparticipation Options

Next, we present the differentiated model with nonparticipation options. This is the model previously studied in Cai and Singham (2018). The model takes the form (6) in terms of the contract options offered.

$$\begin{bmatrix} \theta_1 \\ \vdots \\ \theta_d \\ \vdots \\ \theta_D \end{bmatrix} = \begin{bmatrix} \theta_{11} \cdots \theta_{1n} \cdots \theta_{1N} \\ \vdots \\ \theta_{d1} \cdots \theta_{dn} \cdots \theta_{dN} \\ \vdots \\ \theta_{D1} \cdots \theta_{Dn} \cdots \theta_{DN} \end{bmatrix} \implies \begin{bmatrix} (0,0)\ (q_{11}, t_{11}) \cdots (q_{1n}, t_{1n}) \cdots (q_{1N}, t_{1N}) \\ \vdots \\ (0,0)\ (q_{d1}, t_{d1}) \cdots (q_{dn}, t_{dn}) \cdots (q_{dN}, t_{dN}) \\ \vdots \\ (0,0)\ (q_{D1}, t_{D1}) \cdots (q_{Dn}, t_{Dn}) \cdots (q_{DN}, t_{DN}) \end{bmatrix} \tag{6}$$

Note that each contract has a nonparticipation option $(0, 0)$ which can be chosen in any period after the contracting stage. Because the agent has the option after committing to contract $x_{d'}$ to skip participation at future time periods, his expected utility is:

$$u(x_{d'}, \theta_d) = \sum_{n=1}^{N} \pi_{dn} \max\{\max_{n'}\{v(q_{d'n'}, \theta_{dn}) - t_{d'n'}\}, 0\}. \tag{7}$$

The principal's expected profit if the agent with distribution θ_d selects x_d is thus

$$\Phi_d(x_d, \theta_d) = \sum_{n=1}^{N} \pi_{dn} 1_{\{v(q_{dn}, \theta_{dn}) - t_{dn} \geq 0\}} (t_{dn} - s(q_{dn}, \theta_{dn})), \tag{8}$$

where $1_{\{A\}}$ is the indicator function for event A. The principal's formulation is:

$$
\begin{array}{llllll}
\Phi^{DN} = & \max\limits_{\{q_{dn},t_{dn}\}} & \sum_{d=1}^{D} \mu_d \sum_{n=1}^{N} \pi_{dn} 1_{\{v(q_{dn},\theta_{dn})-t_{dn}\geq 0\}}(t_{dn} - s(q_{dn},\theta_{dn})) & & & \\
 & \text{s.t.} & u(x_d,\theta_d) \geq u(x_{d'},\theta_d) & d,d' \in \mathcal{D}, d \neq d' & (IC_{dd'}) & \\
 & & v(q_{dn},\theta_{dn}) - t_{dn} \geq v(q_{dn'},\theta_{dn}) - t_{dn'} & d \in \mathcal{D}, n,n' \in \mathcal{N}, n \neq n' & (IC_{dnn'}) & (9) \\
 & & q_{dn} \geq 0, t_{dn} \geq 0 & d \in \mathcal{D}, n \in \mathcal{N} & (NN_{dn}). &
\end{array}
$$

The presence of indicator variables in the objective function of (9) requires a mixed-integer nonlinear solver. Next, we present an equivalent alternative formulation which does not require indicator variables, thus making it easier to solve numerically using nonlinear solvers.

Proposition 3.4. Consider formulation Φ^{DN} in (9). At optimality, all the indicator variables in the objective are equal to 1, so that $v(q_{dn},\theta_{dn}) - t_{dn} \geq 0$ for all d,n. This allows (9) to be reformulated as the nonlinear program:

$$
\begin{array}{llllll}
\Phi^{DN} = & \max\limits_{\{q_{dn},t_{dn}\}} & \sum_{d=1}^{D} \mu_d \sum_{n=1}^{N} \pi_{dn}(t_{dn} - s(q_{dn},\theta_{dn})) & & & \\
 & \text{s.t.} & u(x_d,\theta_d) \geq u(x_{d'},\theta_d) & d,d' \in \mathcal{D}, d \neq d' & (IC_{dd'}) & \\
 & & v(q_{dn},\theta_{dn}) - t_{dn} \geq v(q_{dn'},\theta_{dn}) - t_{dn'} & d \in \mathcal{D}, n,n' \in \mathcal{N}, n \neq n' & (IC_{dnn'}) & (10) \\
 & & v(q_{dn},\theta_{dn}) - t_{dn} \geq 0 & d \in \mathcal{D}, n, \in \mathcal{N} & (IR_{dn}) & \\
 & & q_{dn} \geq 0, t_{dn} \geq 0 & d \in \mathcal{D}, n \in \mathcal{N} & (NN_{dn}). &
\end{array}
$$

Proof. Call the mixed-integer formulation (9) MI and the nonlinear formulation (10) NLP. Because the optimal solution to NLP is feasible to MI with the same objective function value, we have $\Phi^{NLP} \leq \Phi^{MI}$. To show the formulations are equivalent, we show $\Phi^{MI} \leq \Phi^{NLP}$. Let $\left[(q_{dn}^{MI}, t_{dn}^{MI})_{d=1,\ldots,D,n=1,\ldots,N}\right]$ be an optimal solution to (9) which delivers Φ^{MI}. If this solution is feasible to NLP, then $\Phi^{NLP} = \Phi^{MI}$. If this solution is infeasible to NLP, construct the following feasible solution to NLP for all d,n:

$$
(q_{dn}^{NLP}, t_{dn}^{NLP}) = \begin{cases} (q_{dn}^{MI}, t_{dn}^{MI}) & \text{if} \quad v(q_{dn}^{MI},\theta_{dn}) - t_{dn}^{MI} \geq 0 \\ (0,0) & \text{if} \quad v(q_{dn}^{MI},\theta_{dn}) - t_{dn}^{MI} < 0. \end{cases} \tag{11}
$$

This solution delivers an optimal objective value $\Phi^{NLP} = \Phi^{MI}$ assuming $s(0, \theta_{dn}) = 0$ from Assumption 3.1, thus $\Phi^{MI} \leq \Phi^{NLP}$ and we have $\Phi^{MI} = \Phi^{NLP}$. $\square$

Note formulation Φ^{DN} does not require IR_d constraints from (5) because they are redundant, as the utility will always be nonnegative due to the nonparticipation options.

3.2 Aggregated Models

When N is large, differentiated models might be too complex to implement as it could be infeasible to execute such a large number of options. It may be beneficial to pool contracts and offer one option targeted towards multiple demand levels. This section formulates the heterogeneous agent distribution problem by offering a single option for each demand distribution without differentiating across the different demand levels. As before, we have one model where the agent must participate at each time period once a contract is selected, and a second model with nonparticipation options.

3.2.1 Aggregated Contracts with Forced Participation

The principal offers the agent a menu of contracts, as shown in the RHS of (12),

$$\begin{bmatrix} \theta_1 \\ \vdots \\ \theta_d \\ \vdots \\ \theta_D \end{bmatrix} = \begin{bmatrix} \theta_{11} \cdots \theta_{1n} \cdots \theta_{1N} \\ \vdots \\ \theta_{d1} \cdots \theta_{dn} \cdots \theta_{dN} \\ \vdots \\ \theta_{D1} \cdots \theta_{Dn} \cdots \theta_{DN} \end{bmatrix} \Longrightarrow \begin{bmatrix} (q_1, t_1) \\ \vdots \\ (q_d, t_d) \\ \vdots \\ (q_D, t_D) \end{bmatrix}. \tag{12}$$

The utility and objective functions are similar to those for the forced participation case for differentiation, except there are fewer decisions variables (q_d, t_d) than (q_{dn}, t_{dn}). The utility function of an agent with distribution θ_d for selecting contract $x_{d'}$ is:

$$u(x_{d'}, \theta_d) = \sum_{n=1}^{N} \pi_{dn} \left(v(q_{d'}, \theta_{dn}) - t_{d'} \right). \tag{13}$$

The principal's expected profit if an agent with distribution θ_d chooses contract x_d is

$$\Phi_d(x_d, \theta_d) = \sum_{n=1}^{N} \pi_{dn}(t_d - s(q_d, \theta_{dn})). \tag{14}$$

We formulate the aggregated contracts with forced participation model as follows:

$$\begin{aligned}
\Phi^{AF} = \quad \max_{\{q_d, t_d\}} \quad & \sum_{d=1}^{D} \mu_d \textstyle\sum_{n=1}^{N} \pi_{dn}(t_d - s(q_d, \theta_{dn})) & & \\
\text{s.t.} \quad & u(x_d, \theta_d) \geq 0 & d \in \mathcal{D} \quad & (IR_d) \\
& u(x_d, \theta_d) \geq u(x_{d'}, \theta_d) & d, d' \in \mathcal{D}, d' \neq d \quad & (IC_{dd'}) \\
& q_d \geq 0, t_d \geq 0 & d \in \mathcal{D} \quad & (NN_d).
\end{aligned} \tag{15}$$

We require IR_d to induce agents to choose a contract. The constraints at the option level (IR_{dn} or $IC_{dnn'}$) are not necessary as there are no options to choose from.

3.2.2 Aggregated Contracts with Nonparticipation Options

In this model, each contract only offers one nonzero option and the agent may opt to skip participation at individual time periods in the future. At each time epoch ($t = 1, 2, \ldots,$), he can either choose not to participate or receive q_d units for a total price t_d. The agent selects a contract $x_d = [(0,0), (q_d, t_d)]$ at time $t = 0$ as in (16),

$$\begin{bmatrix} \theta_1 \\ \vdots \\ \theta_d \\ \vdots \\ \theta_D \end{bmatrix} = \begin{bmatrix} \theta_{11} \cdots \theta_{1n} \cdots \theta_{1N} \\ \vdots \\ \theta_{d1} \cdots \theta_{dn} \cdots \theta_{dN} \\ \vdots \\ \theta_{D1} \cdots \theta_{Dn} \cdots \theta_{DN} \end{bmatrix} \Longrightarrow \begin{bmatrix} (0,0), (q_1, t_1) \\ \vdots \\ (0,0), (q_d, t_d) \\ \vdots \\ (0,0), (q_D, t_D) \end{bmatrix}. \tag{16}$$

The formulation for the aggregated contracts with nonparticipation options model is:

$$
\begin{aligned}
\Phi^{AN} = \quad & \max_{\{q_d, t_d\}} && \sum_{d=1}^{D} \mu_d \textstyle\sum_{n=1}^{N} \pi_{dn} 1_{\{v(q_d,\theta_{dn})-t_d \geq 0\}} (t_d - s(q_d, \theta_{dn})) && && \\
& \text{s.t.} && u(x_d, \theta_d) \geq u(x_{d'}, \theta_d) && d, d' \in \mathcal{D}, d' \neq d && (IC_{dd'}) \\
& && q_d \geq 0, t_d \geq 0 && d \in \mathcal{D} && (NN_d),
\end{aligned}
\tag{17}
$$

where

$$u(x_{d'}, \theta_d) = \sum_{n=1}^{N} \pi_{dn} \max\{v(q_{d'}, \theta_{dn}) - t_{d'}, 0\} \tag{18}$$

represents an agent with demand distribution θ_d selecting contract $x_{d'}$ and

$$\Phi_d(x_d, \theta_d) = \sum_{n=1}^{N} \pi_{dn} 1_{\{v(q_d,\theta_{dn})-t_d \geq 0\}} (t_d - s(q_d, \theta_{dn})) \tag{19}$$

is the principal's expected profit from the agent when he chooses contract x_d. The IR_d constraints from the AF model are redundant because the agents always have nonnegative utility from the nonparticipation option. Moreover, because there is no contract differentiation and nonparticipation is allowed, there are no IR_{dn} constraints which ensure $v(q_d, \theta_{dn}) - t_d \geq 0$. This means that when the agent receives a negative utility from participation in future time periods for some demand level, they will choose the nonparticipation option in those cases.

4 ANALYSIS

This section analyzes the four models (denoted in summary as DF, DN, AF, and AN) to derive general insights. First, we compare and contrast the four models. Table 1 shows the constraints needed for each formulation. We see that all formulations have $IC_{dd'}$ constraints, while forced participation contracts

Table 1: Model Constraints

	Differentiated	Aggregated
Forced	$IC_{dd'}, IR_d, IC_{dnn'}$	$IC_{dd'}, IR_d$
Nonparticipation	$IC_{dd'}, IC_{dnn'}$	$IC_{dd'}$

also require IR_d constraints to induce the agents to choose a contract at $t = 0$. Differentiated contracts require $IC_{dnn'}$ constraints to ensure the agent picks the option that best aligns with his true demand.

Tables 2 and 3 summarize the principal's objective functions and the agent's utility functions for the different models for quick reference in the following analysis.

Table 2: Objective Functions $\Phi_d(x_d, \theta_d)$

	Differentiated	Aggregated
Forced	$\sum_{n=1}^{N} \pi_{dn}(t_{dn} - s(q_{dn}, \theta_{dn}))$	$\sum_{n=1}^{N} \pi_{dn}(t_d - s(q_d, \theta_{dn}))$
Nonparticipation	$\sum_{n=1}^{N} \pi_{dn} 1_{\{v(q_{dn},\theta_{dn})-t_{dn}\geq 0\}}(t_{dn} - s(q_{dn}, \theta_{dn}))$	$\sum_{n=1}^{N} \pi_{dn} 1_{\{v(q_d,\theta_{dn})-t_d\geq 0\}}(t_d - s(q_d, \theta_{dn}))$

Table 3: Utility Functions $u(x_{d'}, \theta_d)$

	Differentiated	Aggregated
Forced	$\sum_{n=1}^{N} \pi_{dn} \max_{n'}\{v(q_{d'n'}, \theta_{dn}) - t_{d'n'}\}$	$\sum_{n=1}^{N} \pi_{dn}(v(q_{d'}, \theta_{dn}) - t_{d'})$
Nonparticipation	$\sum_{n=1}^{N} \pi_{dn} \max\{\max_{n'}\{v(q_{d'n'}, \theta_{dn}) - t_{d'n'}\}, 0\}$	$\sum_{n=1}^{N} \pi_{dn} \max\{v(q_{d'}, \theta_{dn}) - t_{d'}, 0\}$

Next, we derive comparisons between models to determine which contracting options are preferred by the principal. We show that the principal prefers differentiation, and when she differentiates, forced participation gives her a higher expected profit.

Proposition 4.1. For the principal, differentiated models always yield optimal profits greater than or equal to those of aggregated models. That is, for any θ, μ, π, and $v(q_{d'n'}, \theta_{dn})$ and $s(q_{d'n'}, \theta_{dn})$ meeting Assumptions (3.1)-(3.3) we have $\Phi^{AF} \leq \Phi^{DF}$ and $\Phi^{AN} \leq \Phi^{DN}$.

Proof. We first show $\Phi^{AF} \leq \Phi^{DF}$. Let (q_d^{AF}, t_d^{AF}), $d = 1, \ldots, D$ be the optimal solution for Φ^{AF}. Construct a candidate solution to DF with

$$q_{dn} = q_d^{AF} \quad t_{dn} = t_d^{AF}, \qquad d \in \mathcal{D}, n \in \mathcal{N}.$$

In this case, the same option is offered to all demand levels n for a given contract d. In DF, IR_d and $IC_{dd'}$ also hold as they hold for AF. Additionally, for a given θ_{dn} in DF, $v(q_{dn'}, \theta_{dn}) - t_{dn'}$ is the same for all n' so $IC_{dnn'}$ is feasible. Because $\Phi^{AF} = \Phi^{DF}$ for this solution, we can construct a feasible solution to DF with an objective function value at least as good as the optimal Φ^{AF}. Hence, we have $\Phi^{AF} \leq \Phi^{DF}$.

The proof for $\Phi^{AN} \leq \Phi^{DN}$ for the non-participating case is similar, in that we can show that the optimal solution to Φ^{AN} can be used to construct a feasible solution to Φ^{DN} that has all the same options equal to those in AN, and this solution has the same objective value. $\square$

Proposition 4.2. When the principal chooses to differentiate options, forced participation always yields profits greater than or equal to those from offering nonparticipation options. That is, $\Phi^{DN} \leq \Phi^{DF}$.

Proof. As shown in Proposition 3.4, the optimal solution to Φ^{DN} has all $v(q_{dn}, \theta_{dn} - t_{dn}) \geq 0$, so the objective functions and utility functions for DF and DN are the same. Hence, the optimal solution to DN yields a feasible solution to DF with the same objective value so $\Phi^{DN} \leq \Phi^{DF}$. □

We have established that the principal prefers differentiated contracts to aggregated ones, and under differentiation prefers forced to nonparticipation. However, under aggregation, the principal's preference between forced and nonparticipation depends on the underlying demand distribution used, and there is no clear dominance between AF and AN. The intuition for why this could happen comes from an example where two different distributions yield the same AF optimal solution, but one distribution yields $\Phi^{AF} < \Phi^{AN}$ while the other has $\Phi^{AF} > \Phi^{AN}$. Consider $D = 1$ and $N = 3$, so that there are three possible demands $\theta_{11}, \theta_{12}, \theta_{13}$. Suppose there are two possible demand distributions $\pi_{11}, \pi_{12}, \pi_{13}$, both with the same expected demand. Sometimes these two distributions will yield the same AF optimal policy because the principal's objective and constraints consist of expectations. Suppose that both θ_{12} and θ_{13} observe nonnegative information rent $(v(q_d, \theta_{dn}) - t_d)$ from the optimal AF policy, but θ_{11} still must participate when realized despite negative information rent. In the AF model, the principal collects profit from all three demands levels. In the AN model, the principal will no longer receive profit from θ_{11}, and may find it optimal to only target θ_{12} and θ_{13}, possibly obtaining a higher profit from them. However, the different probability weights on the demand levels in the two distributions result in different values of optimal profits Φ^{AN} which could be greater or less than Φ^{AF}.

Thus, we find that forced participation is not always better in the aggregated case because a higher profit may be obtained by targeting high demand levels and allowing lower demand levels nonparticipation options, rather than offering an aggregated contract that induces participation in the contracting stage but forces all demand levels to participate in the future.

5 EXPERIMENTAL INSIGHTS

In the experimental results we explore factors which could affect contract preferences. In particular, we look for insights regarding agent preferences. For any particular problem formulation (choice of D,

$N,\ v(q_{d'n'},\ \theta_{dn}),\ s(q_{d'n'}, \theta_{dn}),\ \theta_{dn},\ \pi_{dn})$ the four formulations can be solved and compared to analyze the resulting contracts and agent utility. This section presents overall insights we observed across these numerous experiments.

We ran over 1,000 different possible configurations of input data for each formulation, and calculated and compared optimal solutions for the four models. For the numerical experiments, we used Pyomo to formulate frameworks DF, DN, and AF (Hart et al., 2011, 2012) and used the nonlinear solver IPOPT (Wächter and Biegler, 2006) to generate solutions. Formulation AF is a nonlinear program and is easy to solve, especially because it only has $2d$ decision variables. Formulation Φ^{DF} contains a **max** function in $u(x_{d'}, \theta_d)$ which typically requires binary variables and a mixed-integer nonlinear solver, but we note that the principal's optimization will attempt to increase the t variables and decrease the q variables under Assumption 3.1. In this situation, auxiliary variables with a nonlinear solver can be used. For DN, we employ the modified formulation in Proposition 3.4 to convert the mixed-integer problem to a nonlinear program. This nonlinear program is much faster to solve than its MINLP counterpart. Model AN has indicator variables in the objective function, so this formulation requires a mixed-integer solver to compute a solution. We employed the solver SCIP (Vigerske and Gleixner, 2018) to solve the MINLP.

We tested a variety of value functions meeting Assumptions 3.1 and 3.2 and report a subset of those results here. Tested functions with results not presented here include $v(q, \theta) = \log(q)(1 - e^{-\theta})$ and $v(q, \theta) = \sqrt{q}\theta$ (Akan et al., 2011). We note that value functions not meeting the assumptions can also be computed and analyzed, for example, functions with the term $\theta q - q^2$ which are concave but not always increasing in q. In all experiments, we assume the principal's cost is a linear function of the consumption quantity, and thus $s(q) = \beta q, \beta > 0$. The units for the expected profit and value function are assumed to be in dollars for all experiments.

5.1 Examples with $D = 2, N = 2$

We conducted many experiments using $D = 2$ and $N = 2$ to see the effect of varying demand distribution parameters. Consider the following distribution for agent demand:

$$\begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} = \begin{bmatrix} 1.0 & 1.6 \\ 1.3 - \delta & 1.3 + \delta \end{bmatrix}, \quad \begin{bmatrix} \mu_1 \\ \mu_2 \end{bmatrix} = \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix}, \quad \begin{bmatrix} \pi_1 \\ \pi_2 \end{bmatrix} = \begin{bmatrix} 0.5 & 0.5 \\ 1 - \sigma & \sigma \end{bmatrix}. \tag{20}$$

By varying δ, we can explore cases where the values of the second demand distribution θ_2 fall within the values of the first demand distribution θ_1 (when $\delta \leq 0.3$), and when its values fall outside of the values of the first distribution (when $\delta > 0.3$). The value of 2δ is the range of the demand distribution θ_2. By varying σ, we see the effect of having a higher probability on the higher demand value of the second distribution. We vary both σ and δ jointly across the values $0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9$, to explore the space of possible values of θ_{dn} and π_{dn}. To reduce the dimensions of parametrization, we set $\mu_1 = \mu_2 = 0.5$ and $\pi_{11} = \pi_{12} = 0.5$, though many experiments not reported here modified these values.

Figure 1 shows the results of one experiment with $v(q, \theta) = \alpha e^{\theta} \sqrt{q}$, where $\alpha > 0$ is a scaling coefficient. The left plot shows the optimal expected profit to the principal for all four models, each varying δ and σ. The right plot shows the agent's expected utility in response to the principal's optimal contracts. The agent's expected utility was calculated (given the principal's optimal contracts are employed) using the expressions in Table 4.

Table 4: Expected Utility Across Agents

	Differentiated	Aggregated
Forced	$\sum_{d=1}^{D} \mu_d \sum_{n=1}^{N} \pi_{dn} \max_{n'} \{v(q_{dn'}, \theta_{dn}) - t_{dn'}\}$	$\sum_{d=1}^{D} \mu_d \sum_{n=1}^{N} \pi_{dn} (v(q_d, \theta_{dn}) - t_d)$
Nonpart.	$\sum_{d=1}^{D} \mu_d \sum_{n=1}^{N} \pi_{dn} \max\{\max_{n'} \{v(q_{dn'}, \theta_{dn}) - t_{dn'}\}, 0\}$	$\sum_{d=1}^{D} \mu_d \sum_{n=1}^{N} \pi_{dn} \max\{v(q_d, \theta_{dn}) - t_d, 0\}$

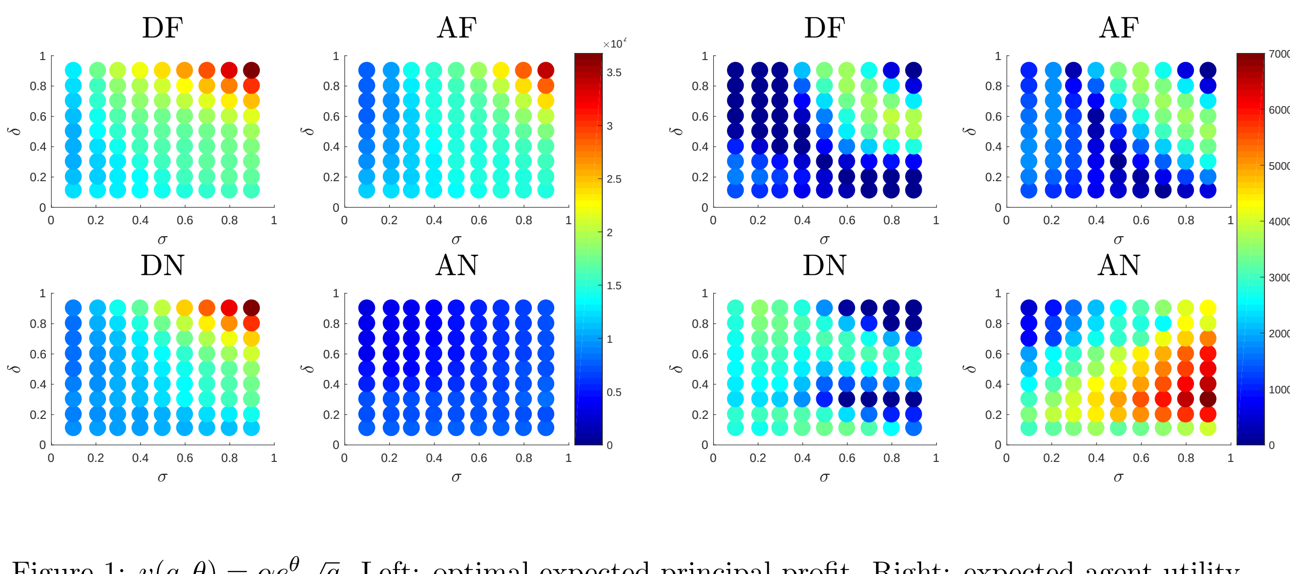

Figure 1: $v(q, \theta) = \alpha e^{\theta} \sqrt{q}$. Left: optimal expected principal profit. Right: expected agent utility.

The left plot of Figure 1 confirms Propositions 4.1 and 4.2. That is, the principal's expected profit is highest in the DF model for all combinations of σ and δ. This is because the DF model allows the principal to design contracts that consider the variability of the demand distributions and thus offers options better aligned with the agents' demand levels. Further, as the probability of the higher demand level (σ) of the agent with demand distribution θ_2 (called a type-2 agent henceforth) increases, the principal's expected profit increases in the DF, AF and AN models. The relationship between the principal's profitability and δ, however, is not as straightforward. When the likelihood of type-2 agents having a high demand level is sufficiently high ($\sigma \geq 0.3$), the principal's expected profit increases in δ in the DF, AF and DN models. On the other hand, if σ is small (≤ 0.2), the principal's expected profit exhibits a non-monotonic pattern in the DF and DN models, while it decreases in the AF model. For the AN model, the principal's expected profit varies little in either δ or σ.

The right plot of Figure 1 shows that neither of the four models generates the highest agent utility for all values of σ and δ. Though the agent, in general, obtains a higher utility in the AN model, there are some exceptions. When $\sigma = 0.5$ or 0.6 and $0.7 \leq \delta \leq 0.9$, both the DF and AF pricing models give the agents a higher aggregated utility. Also, when $\sigma \leq 0.3$ and $\delta \geq 0.6$, the agents prefer the pricing of the DN model. Moreover, the agent's expected utility is monotonic in neither σ nor δ in any of the four models.

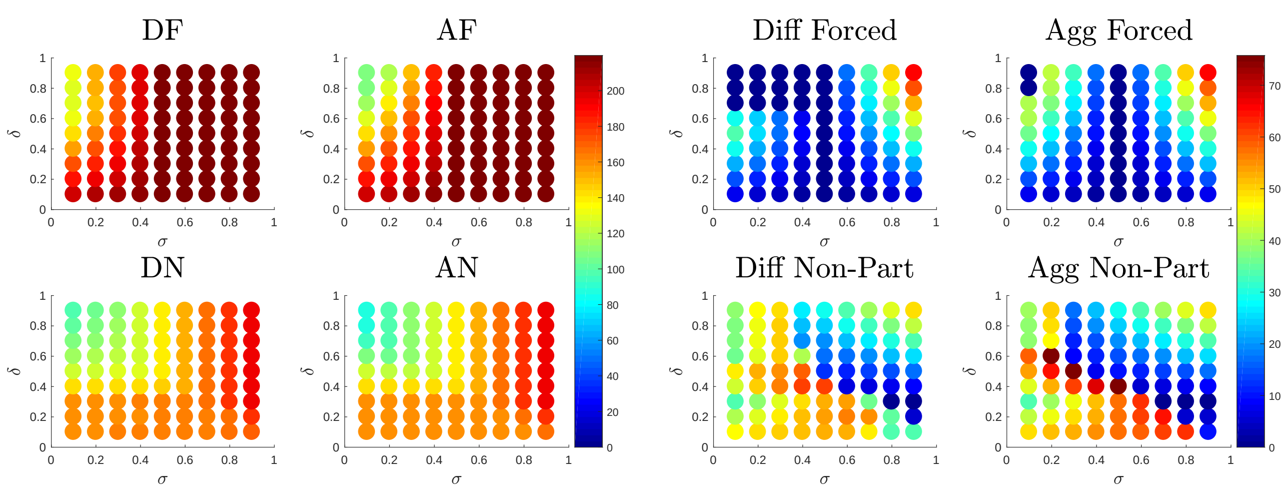

Figure 2: $v(q, \theta) = \alpha\theta(1 - e^{-q})$. Left: optimal expected principal profit. Right: expected agent utility.

Figure 2 gives the results for value function $v(q, \theta) = \alpha\theta(1 - e^{-q})$, $\alpha > 0$. Similar to the results generated from value function $v(q, \theta) = \alpha e^{\theta}\sqrt{q}$, the principal's expected profit increases in σ in all four models. The relationship between the principal's expected profit and δ, however, differs from what is shown in Figure 1. The principal's expected profit decreases in δ when the probability of type-2 agents having a large demand is small ($\sigma \leq 0.4$ in the DF and AF models, $\sigma \leq 0.6$ in the DN and AN models), while it increases in δ when σ is sufficiently large ($\sigma \geq 0.5$ in the DF and AF models, $\sigma \leq 0.7$ in the DN and AN models). For most of the values of σ and δ, the AN model provides a higher expected utility for the agent. Nonetheless, when both σ and δ are close to 0.9 (top right corner), both the DF and AF models give a higher aggregated utility to the agents.

The key point from these figures is that while there may be some general trends across models and parameters, only a computational solution will reveal which situation is best for the agent. Even in the case where demand distributions are simple with $D = 2$ and $N = 2$, the solutions can vary significantly and depend on the exact values used in the demand distribution. Changing the value functions also significantly affects the gradients in the figures.

5.2 An Example with $D = 3, N = 6$

We consider a problem with larger dimensions to model a realistic example, where we choose $D = 3, N = 6$. The example is based off a meal delivery service where customers have varying demand distributions for the number of meals they want to be delivered each week. The principal, who supplies the meals, has a fixed cost per meal, say $s(q, \theta) = \beta q$ where $\beta =$\$6. The agent has the following demand distributions:

$$\begin{bmatrix} \theta_1 \\ \theta_2 \\ \theta_3 \end{bmatrix} = \begin{bmatrix} 2 & 3 & 4 & 5 & 6 & 12 \\ 4 & 8 & 10 & 12 & 14 & 16 \\ 10 & 12 & 14 & 16 & 18 & 24 \end{bmatrix}, \qquad \begin{bmatrix} \mu_1 \\ \mu_2 \\ \mu_3 \end{bmatrix} = \begin{bmatrix} 0.25 \\ 0.25 \\ 0.50 \end{bmatrix}, \tag{21}$$

$$\begin{bmatrix} \pi_1 \\ \pi_2 \\ \pi_3 \end{bmatrix} = \begin{bmatrix} 1/16 & 1/16 & 1/4 & 1/8 & 1/4 & 1/4 \\ 1/6 & 1/6 & 1/6 & 1/6 & 1/6 & 1/6 \\ 1/12 & 1/12 & 1/12 & 1/4 & 1/4 & 1/4 \end{bmatrix}. \tag{22}$$

In this case, the agent can either have a low (θ_1), medium (θ_2), or high (θ_3) demand distribution. Half of the agents follow the high demand distribution ($\mu_3 = 0.5$) with the remaining agents split between the low and medium distributions ($\mu_1 = \mu_2 = 0.25$). The agent has a value function $v(q,\theta) = \alpha\theta(1 - e^{-\gamma\frac{q}{\theta}})$ where $\alpha = 10, \gamma = 2$. When the agent receives zero quantity, his valuation is zero. His valuation increases in the proportion of demand received, q/θ. Any increase in quantity beyond θ results in little increase in his valuation. As in the other value functions considered, the single-crossing property is satisfied as $\frac{\partial^2 v(q,\theta)}{\partial q \partial \theta} = \frac{\partial}{\partial \theta}\left(\frac{\partial v}{\partial q}\right) = \frac{\partial}{\partial \theta}\left(\alpha\gamma \cdot e^{-\gamma\frac{q}{\theta}}\right) = \alpha\gamma \cdot e^{-\gamma\frac{q}{\theta}} \cdot \frac{\gamma q}{\theta^2} > 0$.

We obtain the following optimal contracts for each of the four models in (23) - (25). Instead of reporting the total price, we report the price per meal delivered $\bar{t} = t/q$ which is more meaningful. For each contract option, we report $(q, \bar{t})$. In reality, the principal would need to deliver integer quantities, but we report the true optimal solution here for analysis:

$$[(q^{DF}, \bar{t}^{DF})] = \begin{bmatrix} (0.4, 19.7) & (0.5, 19.0) & (0.9, 17.0) & (0.9, 17.0) & (0.9, 17.0) & (3.4, 14.9) \\ (0.9, 18.8) & (2.2, 15.7) & (2.8, 15.0) & (3.3, 14.6) & (3.4, 14.5) & (3.4, 14.5) \\ (6.0, 14.8) & (7.2, 13.3) & (8.4, 12.3) & (9.6, 11.5) & (10.8, 10.9) & (14.4, 9.6) \end{bmatrix}, \tag{23}$$

$$[(q^{DN}, \bar{t}^{DN})] = \begin{bmatrix} (0, 0) & (0.1, 19.3) & (0.2, 18.8) & (0.9, 16.0) & (0.9, 16.0) & (0.9, 16.0) & (5.1, 13.0) \\ (0, 0) & (0.4, 18.1) & (1.9, 15.7) & (2.7, 14.8) & (4.1, 13.6) & (5.3, 12.9) & (5.3, 12.9) \\ (0, 0) & (2.4, 15.9) & (3.4, 14.9) & (4.5, 14.0) & (8.4, 11.7) & (9.2, 11.3) & (14.4, 9.9) \end{bmatrix}, \tag{24}$$

$$[(q^{AF}, \bar{t}^{AF})] = \begin{bmatrix} (1.1, 16.4) \\ (2.4, 14.9) \\ (10.3, 11.0) \end{bmatrix}, \qquad [(q^{AN}, \bar{t}^{AN})] = \begin{bmatrix} (0, 0) & (1.7, 13.6) \\ (0, 0) & (3.7, 13.1) \\ (0, 0) & (9.1, 11.2) \end{bmatrix}. \tag{25}$$

Furthermore, (26) lists the agent's expected utility by demand distribution where Ψ^{DF} corresponds to the DF model, and so on. Equation (27) reports the principal's expected profit for each model.

$$\Psi^{DF} = \begin{bmatrix} 0.0 \\ 1.8 \\ 6.5 \end{bmatrix}, \qquad \Psi^{DN} = \begin{bmatrix} 1.0 \\ 3.5 \\ 10.7 \end{bmatrix}, \qquad \Psi^{AF} = \begin{bmatrix} 0.0 \\ 1.5 \\ 5.8 \end{bmatrix}, \qquad \Psi^{AN} = \begin{bmatrix} 2.5 \\ 5.2 \\ 11.3 \end{bmatrix}, \tag{26}$$

$$\Phi^{DF} = \$36.0, \qquad \Phi^{DN} = \$33.0, \qquad \Phi^{AF} = \$33.6, \qquad \Phi^{AN} = \$27.9. \tag{27}$$

First, we compare the optimal contracts between the DF and DN models. For the low and medium demand distributions (θ_1 and θ_2), while the DF model offers higher quantities for the lower demand levels, it provides lower quantities for the higher demand levels. However, the DF model assigns higher quantities for all demand levels to the high demand distribution (θ_3). As a result, the overall expected consumption quantity using the pricing of the DF model (6.3 meals) is higher than that of the DN model (5.7 meals). The unit price for each demand distribution decreases as the demand level increases for both the DF and DN models, giving a quantity discount appeal to the agents. Because all three types of agents receive lower utility in the DF model, the principal achieves a higher expected profit in the DF model.

Next, when compared to the AN model, the AF model assigns lower quantities while charging higher prices for both the low and medium demand distributions (θ_1 and θ_2). For the high demand distribution (θ_3), the AF model provides a higher quantity at a slightly lower price. The principal can achieve a higher profit in the AF model as the agents' utilities are lower.

Finally, we compare the aggregated models against the differentiated models. In order to incentivize the agents to participate, the quantities offered in the aggregated models are closer to the lower demand levels of θ_1 and θ_2. For the high demand distribution (θ_3), however, the quantity offered in the aggregated models are designed toward higher demand levels. The principal's profitability is higher in the DN model than the AN model as the agents receive lower utility in the former model. Interestingly, even though the agents receive *higher* utility in the DF model than the AF model, the expected profit to the principal is higher in the DF model because the overall consumption quantity is higher in the DF model (expected meals consumed is 6.3) than the AF model (expected meals consumed is 6.0). This further demonstrates the superiority of the DF model as it not only provides the highest profit to the principal, but also gives the agents higher expected utility in certain cases.

For models where D and N are large, the computational results can reveal significant information about the best type of contract for the principal and the agent. We note that the computing time was minimal for each of these models, with DF, DN, and AF solving nearly instantaneously on a single processor, and with AN taking under a minute.

6 CONCLUSIONS

In this paper, we develop pricing models for services to heterogeneous agents with stochastic demand under the principal-agent framework. Both *ex ante* and *ex post* information asymmetry are considered. The agent can have one of multiple possible demand distributions, and the different values of the demand distribution are realized in each time period after contracting has occurred. We present a four-model framework for organizing different flexible contracting options and employ a computational scheme to deliver optimal contracts to the principal.

We find that the principal prefers to provide differentiated contracts with multiple options designed for different demand levels rather than aggregating different demand levels into one contract option. When differentiation is used, the principal prefers to force participation at each time period after the initial contracting stage, rather than allowing for nonparticipation options. Once the principal's optimal contracts are derived, we can calculate the expected utility across the different agents. The agents generally prefer the aggregated contracts with nonparticipation options as these types of contracts grant the agent more flexibility and a higher expected utility. However, we have identified certain combinations of demand distributions where the agent can receive a higher utility in either differentiated contracts or aggregated contracts with forced participation.

ACKNOWLEDGEMENTS

We are grateful to the National Science Foundation for grants CMMI-1535762 and CMMI-1535831 which supported this work.

References

Akan, Mustafa, Bariş Ata, Martin A. Lariviere. 2011. Asymmetric information and economies-of-scale in service contracting. *Manufacturing & Service Operations Management* **13**(1) 58–72.

Arya, Anil, Brian Mittendorf. 2006. Project assignments when budget padding taints resource allocation. *Management Science* **52**(9) 1345–1358.

Babich, Volodymyr, Hantao Li, Peter Ritchken, Yunzeng Wang. 2012. Contracting with Asymmetric Demand Information in Supply Chains. *European Journal of Operational Research* **217**(2) 333–341.

Bertsimas, Dimitris, Aurélie Thiele. 2006. A robust optimization approach to inventory theory. *Operations Research* **54**(1) 150–168.

Bhargava, Hemant K., Vidyanand Choudhary. 2008. When is versioning optimal for information goods? *Management Science* **54**(5) 1029–1035.

Burnetas, Apostolos, Stephen M Gilbert, Craig E Smith. 2007. Quantity discounts in single-period supply contracts with asymmetric demand information. *IIE Transactions* **39**(5) 465–479.

Cai, Wenbo, Dashi I. Singham. 2018. Principal-agent problems with heterogeneous agent demand distributions. *European Journal of Operational Research* **264**(1) 239–256.

Çakanyildirim, Metin, Qi Feng, Xianghua Gan, Suresh P. Sethi. 2012. Contracting and coordination under asymmetric production cost information. *Production and Operations Management* **21**(2) 345–360.

Cecchini, Mark, Joseph Ecker, Michael Kupferschmid, Robert Leitch. 2013. Solving nonlinear principal-agent problems using bilevel programming. *European Journal of Operational Research* **230**(2) 364–373.

Chen, Fangruo, Guoming Lai, Wenqiang Xiao. 2016. Provision of incentives for information acquisition: Forecast-based contracts vs. menus of linear contracts. *Management Science* **62**(7) 1899–1914.

Chen, YingJu, Mingcherng Deng. 2015. Information sharing in a manufacturer–supplier relationship: Suppliers' incentive and production efficiency. *Production and Operations Management* **24**(4) 619–633.

Corbett, Charles J., Deming Zhou, Christopher S. Tang. 2004. Designing supply contracts: Contract type and information asymmetry. *Management Science* **50**(4) 550–559.

Dempe, Stephan. 1995. Computing optimal incentives via bilevel programming. *Optimization* **33**(1) 29–42.

Ehtamo, Harri, Kimmo Berg, Mitri Kitti. 2010. An adjustment scheme for nonlinear pricing problem with two buyers. *European Journal of Operational Research* **201**(1) 259–266.

Hart, W.E, C. Laird, J-P. Watson, D.L Woodruff. 2012. *Pyomo - Optimization modeling in Python*, vol. 67. Springer Science & Business Media.

Hart, W.E, J-P. Watson, D.L. Woodruff. 2011. Pyomo: Modeling and solving mathematical programs in python. *Mathematical Programming Computation* **3**(3) 219–260.

Iyer, Ananth V., Leroy B. Schwarz, Stefanos A. Zenios. 2005. A principal-agent model for product specification and production. *Management Science* **51** 106–119.

Jiang, Houyuan, Zhan Pang, Sergei Savin. 2012. Performance-based contracts for outpatient medical services. *Manufacturing & Service Operations Management* **14**(4) 654–669.

Kim, Sang-Hyun, Serguei Netessine. 2013. Collaborative cost reduction and component procurement under information asymmetry. *Management Science* **59**(1) 189–206.

Laffont, Jean-Jacques, David Martimort. 2009. *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press.

Lambert, Richard A. 2006. Agency theory and management accounting. *Handbooks of Management Accounting Research* **1** 247–268.

Löffler, Clemens, Thomas Pfeiffer, Georg Schneider. 2012. Controlling for supplier switching in the presence of real options and asymmetric information. *European Journal of Operational Research* **223**(3) 690–700.

Maskin, E., J. Riley. 1984. Monopoly with incomplete information. *RAND Journal of Economics* **15**(2) 171–196.

Özer, Özalp, Wei Wei. 2006. Strategic commitment for optimal capacity decision under asymmetric forecast information. *Management Science* **52**(8) 1238–1257.

Perakis, Georgia, Guillaume Roels. 2010. Robust controls for network revenue management. *Manufacturing & Service Operations Management* **12**(1) 56–76.

Singham, Dashi I. 2018. Sample average approximations for the continuous type principal-agent problem. *Under Review.*

Sundararajan, Arun. 2004. Nonlinear pricing of information goods. *Management Science* **50**(12) 1660–1673.

Taylor, Terry A., Wenqiang Xiao. 2009. Incentives for retailer forescasting: Rebates vs. returns. *Management Science* **55**(10) 1654–1669.

Tran, Thanh, Ramarao Desiraju. 2017. Group-buying and channel coordination under asymmetric information. *European Journal of Operational Research* **256**(1) 68–75.

Vigerske, Stefan, Ambros Gleixner. 2018. SCIP: Global optimization of mixed-integer nonlinear programs in a branch-and-cut framework. *Optimization Methods and Software* **33**(3) 563–593.

Wächter, A., L.T. Biegler. 2006. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical programming* **106**(1) 25–57.

Yang, Zhibin, Göker Aydn, Volodymyr Babich, Damian R. Beil. 2009. Supply disruptions, asymmetric information, and a backup production option. *Management Science* **55**(2) 192–209.

Yang, Zhibin, Göker Aydn, Volodymyr Babich, Damian R. Beil. 2012. Using a dual-sourcing option in the presence of asymmetric information about supplier reliability: Competition vs. diversification. *Manufacturing & Service Operations Management* **14**(2) 202–217.

Zhang, Hao, Mahesh Nagarajan, Greys Sošić. 2010. Dynamic supplier contracts under asymmetric inventory information. *Operations Research* **58**(5) 1380–1397.