Stochastic Mechanism Design
(Extended Abstract)*

Samuel Ieong[1], Anthony Man-Cho So[2], and Mukund Sundararajan[1]
{sieong,manchoso,mukunds}@cs.stanford.edu

[1] Department of Computer Science, Stanford University
[2] Department of Sys. Eng. & Eng. Mgmt., The Chinese University of Hong Kong

Abstract. We study the problem of welfare maximization in a novel setting motivated by the standard stochastic two-stage optimization with recourse model. We identify and address algorithmic and game-theoretic challenges that arise from this framework. In contrast, prior work in algorithmic mechanism design has focused almost exclusively on optimization problems without uncertainty. We make two kinds of contributions.
First, we introduce a family of mechanisms that induce truth-telling in general two-stage stochastic settings. These mechanisms are *not* simple extensions of VCG mechanisms, as the latter do not readily address incentive issues in multi-stage settings. Our mechanisms implement the welfare maximizer in *sequential ex post* equilibrium for risk-neutral agents. We provide formal evidence that this is the strongest implementation one can expect.
Next, we investigate algorithmic issues by studying a novel combinatorial optimization problem called the *Coverage Cost* problem, which includes the well-studied Fixed-Tree Multicast problem as a special case. We note that even simple instances of the stochastic variant of this problem are #P-Hard. We propose an algorithm that approximates optimal welfare with high probability, using a combination of sampling and supermodular set function maximization—the techniques may be of independent interest. To the best of our knowledge, our work is the first to address both game-theoretic *and* algorithmic challenges of mechanism design in multi-stage settings with data uncertainty.

1 Introduction

Welfare maximization has been a central problem in both computer science and economics research. Much work to-date, especially in algorithmic mechanism design, has focused on welfare maximization in deterministic settings [11, 14, 15]. In this paper, we identify and address new challenges that arise in stochastic optimization frameworks. In particular, we consider both algorithmic and incentive issues motivated by the *two-stage stochastic optimization with recourse*

* A journal version of the paper is under preparation and a draft can be found at the authors' websites. Samuel Ieong is funded by a Stanford Graduate Fellowship and NSF ITR-0205633.

model, a model that has been studied extensively in both the operations research community [5], and the computer science community [9, 16, 17].

Roughly speaking, two-stage optimization requires a decision maker (the *center*) to make sequential decisions. In the first stage, given a probability distribution over possible problem instances (called *scenarios*), the center deploys some resources and incurs some cost. Typically, such an initial deployment is not a feasible solution to every possible scenario, but represents a *hedge* on the center's part. In the second stage, once a specific scenario is realized, the center may take *recourse* actions to augment its initial solution to ensure feasibility, and incurs an additional cost for doing so. The goal of the center is to minimize its *expected* cost (or maximize its expected profit).

In this paper, we are interested in situations where the uncertainty is initially unknown to the center, and it needs to learn this information from *selfish agents* in order to maximize social welfare. However, an agent may lie about its private information to improve its utility. To solve this informational problem, the center has to interleave elicitation and optimization. In order to appreciate the challenges that arise from the stochastic setting, let us first consider a stochastic variant of the well-studied *Fixed Tree Multicast (FTM)* problem [1, 6, 13].

Recall that an FTM instance consists of a tree T with undirected edges and a designated node called the *root*. A set of players, $U = \{1, 2, \ldots, n\}$, are located at the nodes of the tree. Each player $i \in U$ is interested in a service provided by the root and has a private value θ_i for being *served*. Serving a user involves building the path from the root to the node at which the user is located. The center serves a set $S \subseteq U$ of users by building edges in the union of paths that correspond to the serviced nodes, and pays the costs of the edges built.

In our stochastic two-stage formulation, there is initially some uncertainty regarding the values of players being served. This uncertainty is modeled as a distribution over values for each player, and is resolved in the second stage when each learns of its value. Both the distribution and the value are private to the player. Edges can be built in either the first or the second stage, with the costs being higher in the second stage for the corresponding edge. Such an increase in costs can be viewed as a premium for the extra information obtained in the second stage. A precise formulation is given in Section 4.

Our objective is to maximize expected social welfare — the sum of the values of the players served less the cost incurred. *What are the challenges introduced by the two-stage stochastic setting?*

The first challenge is game-theoretic. While Vickrey-Clarke-Groves (VCG) mechanisms can induce players to report their true information in single-shot settings [8], they do not apply directly to the two-stage setting. We demonstrate that it is possible to induce truth-telling behavior in *sequential ex post* equilibrium via an explicit construction of a two-stage mechanism. This solution concept is different from classical ex post implementation, and will be further explained in Section 4. We also formally argue that this is the strongest implementation one can expect, by showing that it is impossible to construct a mechanism that implements the social objective in dominant strategies.

The second challenge is algorithmic. We consider a novel combinatorial optimization problem called the *Coverage Cost* problem (see Section 5 for the precise formulation) to investigate algorithmic issues that arise in such settings. The Coverage Cost (CC) problem contains FTM as a special case. We find that maximizing welfare can be difficult even when the deterministic version is easy to solve. For instance, maximizing welfare in a deterministic, single-shot version of FTM can be solved by a linear time algorithm [6]. On the other hand, maximizing welfare for a stochastic version of FTM is #P-hard (Theorem 5). We then develop an algorithm for stochastic CC problems that yields an additive approximation to the optimal expected welfare with high probability. Our solution is based on a combination of sampling techniques (see, e.g., [12]) and supermodular function maximization [10] (Theorem 6), and may be of independent interest.

Due to space constraints, most proofs have been omitted in this extended abstract. Readers interested in more details can find the proofs in the full paper.

2 Related Work

A few recent papers have focused on *dynamic mechanisms*, under the setting of Markov Decision Processes [2, 3]. Our work differs from these in three respects. First, to the best of our knowledge, this is the first time an algorithmic aspect of a two-stage mechanism has been studied. The computational hardness leads to the use of a sampling-based approximation algorithm, and we describe the precise trade-off between incentive compatibility and computational efficiency (Theorem 4). Second, by application of backward induction, we identify a *family* of incentive compatible mechanisms, rather than a single mechanism. Finally, we introduce the *sequential* generalization of classical solution concepts, and formally argue why a stronger incentive guarantee — namely an implementation in dominant strategies — is impossible to achieve (Theorem 3). We expect our impossibility result to be applicable to the works mentioned above.

To the best of our knowledge, our algorithmic result, i.e. approximating maximum welfare for a stochastic coverage cost instance, is not obtainable via current techniques. For instance, the technique in [4] requires the objective function to be non-negative for all possible actions and scenarios. This condition does not hold in our problem. The technique in [12] does not address how the underlying problem is to be solved, and yields a different bound that depends on the variance of a certain quantity.

3 Stochastic Welfare Maximization

In this section, we define stochastic welfare maximization in a general setting. The terminology and the general definition introduced in this section are motivated by mechanism design. We review two-stage stochastic optimization in the appendix and refer the interested readers to the survey [17].

Informally, in a two-stage stochastic welfare maximization problem, the center decides on the eventual social outcome in two stages. In the first stage, with less information available, the center commits some resources, at a cost. In the second stage, with additional, precise information available, it augments initial allocation by performing (typically more expensive) *recourse* actions.

Formally, in the first stage the center picks an outcome from the set of feasible outcomes $\mathcal{O}^1$, incurring a first-stage cost $c^1 : \mathcal{O}^1 \mapsto \mathbb{R}$. In the second stage, the center may augment the first stage allocation by picking an outcome from the set $\mathcal{O}^2$, incurring an additional cost $c^2 : \mathcal{O}^1 \times \mathcal{O}^2 \mapsto \mathbb{R}$. Note that the second-stage cost depends on both the first and second-stage choices.

Next, we describe the relationship between agent types and their valuations.[3] Let Θ_i be the (ground) type space of agent i, for $i = 1, \ldots, n$. Let $v_i : \Theta_i \times \mathcal{O}^1 \times \mathcal{O}^2 \mapsto \mathbb{R}$ denote i's *valuation*. In other words, an agent's valuation depends on its realized type and the outcomes of both stages.

The ground type of an agent is revealed in two stages. In the first stage, an agent i only learns of a *probability distribution* δ_i over its ground types. We call this distribution the agent's *supertype*, and denote the supertype space of agent i by Δ_i. Its elements, $\delta_i \in \Delta_i$, are distributions on Θ_i. In the second stage, agent i learns of its *ground type* (or *type* for short), realized according to the distribution δ_i that is *independent* of other agents' type realizations. We call the collective realized types of all agents in the system a *scenario*, corresponding to a scenario in two-stage stochastic optimization.

Most work on two-stage stochastic optimization focus on minimizing cost. In contrast, we are interested in maximizing *social welfare* as defined below:

Definition 1 The social welfare of outcomes $x^1 \in \mathcal{O}^1$ and $x^2 \in \mathcal{O}^2$ in scenario $\theta = (\theta_1, \ldots, \theta_n)$ is:

$$SW(\theta, x^1, x^2) = \sum_{i=1}^{n} v_i(\theta_i, x^1, x^2) - c^1(x^1) - c^2(x^1, x^2) \tag{1}$$

As first-stage outcomes are picked without precise information on agent types, we focus on maximizing *expected* social welfare, i.e.

$$\max \mathbb{E}_{\theta \sim \delta} \left[SW(\theta, x^1, x^2) \right] \tag{2}$$

where $\theta \sim \delta$ means that the scenario vector θ is distributed according to δ.

4 Mechanism Design Formulation

We now address the first challenge in stochastic welfare maximization, that of eliciting the supertypes and the realized types from selfish agents. Our treatment is fully general, and applies to any two-stage stochastic optimization problems.

First, let us define agents' utility functions. We assume that agents have *quasi-linear* utilities. If t is the transfer to agent i, then he has utility:

$$u_i(\theta_i, x^1, x^2, t) = v_i(\theta_i, x^1, x^2) + t \tag{3}$$

[3] For notation, when a type/supertype (space) is subscripted, it refers to that of a particular agent; when it is not, it refers to the Cartesian product over the agents.

We also assume that agents are *risk-neutral* in the first stage, i.e., they look to maximize their expected utility over the distribution of scenarios.

The mechanism design framework is as follows:

Definition 2 A *two-stage stochastic mechanism* is parametrized by a pair of mechanisms, $(\langle f^1, \{t_i^1\}_{i=1}^n\rangle, \langle f^2, \{t_i^2\}_{i=1}^n\rangle)$, where:

1. Initially, each agent i has a supertype $\delta_i \in \Delta_i$. The first-stage mechanism accepts "supertype" bids from agents.
2. The mechanism applies the decision rule, $f^1 : \Delta \mapsto \mathcal{O}^1$ to pick a first-stage outcome as a function of declared supertypes. It applies the transfer functions $t_i^1 : \Delta \mapsto \mathbb{R}$ to determine first-stage transfers for each agent i.
3. Each agent i now realizes its type θ_i according to the distribution specified by the supertype δ_i. The mechanism accepts "type" bids from each agent.
4. The second-stage mechanism applies the decision rule, $f^2 : \Delta \times \mathcal{O}^1 \times \Theta \mapsto \mathcal{O}^2$ and picks a second-stage outcome as a function of the declared types, the declared supertypes, and the first-stage outcome. It applies the transfer functions $t_i^2 : \Delta \times \mathcal{O}^1 \times \Theta \mapsto \mathbb{R}$ to determine second-stage transfers for each agent i. At this stage the utility of each agent i, $u_i(\theta_i, x^1, x^2, t_i^1 + t_i^2)$, is determined based on its true type, the two outcomes, and the two transfers.

We model the game induced by the two-stage stochastic mechanism among the agents as a *dynamic game of incomplete information.* The strategy of each agent specifies its actions for each of its information sets. Note that an agent's second-stage action may depend on the first stage outcome, the agent's supertype and the agent's realized type. Thus, we define the strategy of agent i with supertype δ_i and type θ_i to be $s_i(\delta_i, \theta_i) = \langle s_i^1(\delta_i), s_i^2(\delta_i, x^1, \theta_i)\rangle$, where x^1 is the (publicly observable) first-stage decision made by the center, and s_i^1 and s_i^2 are the strategy mappings of agent i in the two stages.

4.1 *Sequential* Solution Concepts

Before explaining our mechanism, let us first consider what solution concept is appropriate for our setting. Classical solution concepts, including dominant-strategy (DS), ex post (EP), and Bayes-Nash (BN) equilibrium, all focus on whether an agent has incentive to deviate from truth-telling *knowing its own type.* For example, in the classical BN equilibrium, an agent cannot deviate from its strategy and improve its expected utility, where the expectation is taken over the distribution of the other agent's types. In contrast, in our two-stage setting, an agent is also uncertain about its own realized type in the first stage. The uncertainty about an agent's own type makes these classical concepts inappropriate for our setting. Formally,

Theorem 1 *For general two-stage stochastic optimization problems, if the outcomes picked by the mechanism depend on both the supertype and ground type of an agent, then there exists a supertype space for which the agent may have incentive to lie about its supertype if he foresees its realized type. This holds even if the agent is the only participant in the mechanism.*

The impossibility result is based on a public-good problem where the center can decide to serve the agent in either the first stage, the second stage, or not at all, and the agents may have a high or low type realization.

In order to match the flow of the information in the execution of the mechanism with the timing of agents' reports, we introduce a *sequential* generalization of the classical solution concepts. Informally, a sequential solution is one where agents have no incentive to deviate from their equilibrium strategy given the information available up to the time they take an action. Applied to our setting, a set of strategies is in sequential EP equilibrium if an agent

- cannot improve its *expected utility* by lying in the first stage, where the expectation is taken over the scenarios, even if he knows the other agents' true supertype, provided the other agents are truthful; and
- cannot improve its utility by lying in the second stage[4].

4.2 A *Sequential Ex Post* Implementation of Welfare Maximizer

Since we are interested in implementing the welfare maximizer, the decision rules for both stages are fixed, and our goal is to find transfer functions such that truth-telling by all agents constitutes a sequential EP equilibrium.

We start by noting that once first-stage decisions have been made, the situation resembles a standard one-shot VCG setting. Hence, we have:

Lemma 1 *For any first-stage decisions $\overline{x}^1 \in \mathcal{O}^1$, first-stage payments t^1, realization of types θ, the family of* Groves mechanism *implements the social welfare maximizer, conditional on the first-stage decisions, in dominant strategies.*

Henceforth, we set the second-stage transfer function to be:

$$t_i^{2*}(\delta, \overline{x}^1, \hat{\theta}) = \sum_{j \neq i} v(\hat{\theta}_j, \overline{x}^1, x^{2*}) - c^2(\overline{x}^1, x^{2*}) + g_i^2(\delta_{-i}, \theta_{-i}) + h_i^2(\delta, \hat{\theta}_{-i})$$

where $g_i^2(\cdot, \cdot)$ is an arbitrary function that does not dependent on either δ_i or θ_i, and $h_i^2(\cdot, \cdot)$ is an arbitrary function that does not depend on θ_i.

We next apply the technique of *backward induction* to analyze the first stage of the dynamic game. Suppose that we fix our second-stage mechanism to be a Groves mechanism $\langle f^{2*}, \{t_i^{2*}\}_{i=1}^n \rangle$. When we evaluate the expected utility of an agent's first-stage strategy, we can assume that all agents will truthfully report their second-stage realized types. By propagating the expected transfers in the second stage to the first stage, we find the following family of transfer functions that helps to implement the first-stage decision rule truthfully. The proof can be found in the full paper.

Theorem 2 *Let x^{1*} be the optimal first-stage decisions based on the declared supertypes $\hat{\delta}$. Let the first-stage transfers be given by:*

$$t_i^{1*}(\hat{\delta}) = -c^1(x^{1*}) + h_i^1(\hat{\delta}_{-i}) - \mathbb{E}_{\theta_{-i} \sim \delta_{-i}}[h_i^2(\hat{\delta}, \theta_{-i})]$$

[4] In fact, in our mechanism, truth-telling is weakly dominant in the second stage.

where $h_i^1(\cdot)$ is an arbitrary function that does not depend on the declaration $\hat{\delta}_i$ of agent i. Then, together with any Groves mechanism in the second stage, the two-stage mechanism implements the expected social welfare maximizer in sequential ex post *equilibrium.*

For a concrete application of this theorem, we consider the problem of implementing the social welfare maximizer for a class of problems known as the stochastic two-stage *coverage cost* problems in Section 5, which includes (stochastic) public goods and FTM problems as special cases.

Similar mechanisms have been proposed in [2, 3]. Our results differ in that our proof is based on an explicit backward induction analysis. As a result, we obtain a family of incentive compatible mechanisms, of which the mechanisms in [2, 3], when specialized to a two-stage setting, are members of the family.

4.3 Impossibility of Sequential Dominant Strategy Implementation

A stronger form of incentive compatibility than sequential EP equilibrium is that of sequential DS equilibrium. This asserts that truth-telling is a weakly dominant strategy regardless of the other agents' strategies, provided that an agent does not know the future realization. We now show that under mild restrictions on the transfer functions, no mechanism can achieve welfare maximization in DS.

Definition 3 A mechanism satisfies *No Positive Transfers* (NPT) if for all players i, the first and second-stage payments t_i^1, t_i^2 are non-positive.

Definition 4 A mechanism satisfies *Voluntary Participation* (VP) if all truthful players are guaranteed non-negative expected utility and non-negative marginal second-stage utility.

The definition of NPT asserts that all payments flow from the players to the mechanism. The VP condition requires that it is in the agents' interest to participate in the mechanism in both stages. We now state our main theorem.

Theorem 3 *There exists an instance of the two-stage stochastic public goods problem with two players for which no mechanism satisfying VP, NPT can implement the expected welfare maximizer (WLF) in DS.*

Thus, subject to the conditions of NPT and VP, we have shown that our implementation in the previous section is the strongest possible.

Informally, one cannot implement the socially efficient outcome in dominant strategies because when certain agents in the mechanism lie *inconsistently* — for example, by first declaring a "low" distribution in the first stage, followed by a "high" valuation in the second stage — other agents may benefit from misrepresenting their distributions. We now formalize this intuition.

Consider an instance of a two-stage stochastic public goods problem with two players, A and B, with some distributions δ_A, δ_B on their respective values of being served by a public good e. The cost of the public good is $c^1 \gg 0$ in the first stage and $c^2 = 2c^1$ in the second stage. Let h be some value $> c^2$. We now define distributions that play a role in the proof. Let $\hat{H}$ be a degenerate distribution localized at h, H denote a full-support distribution[5] with most of its mass at h, L denote a full-support distribution with most of its mass at 0, and $\hat{M}$ denote the degenerate distribution localized at $c^1/2$.

For notation, let $\langle D^1, D^2 \rangle$ denote the strategy of a player that reports D^1 as its supertype and reports $v \sim D^2$ as its type. When we consider only the first-stage strategy, we may write $\langle D, \cdot \rangle$ instead. A strategy is consistent if it is of the form $\langle D, D \rangle$, and truthful if it is consistent and $D = \delta_i$ for agent i.

The following lemmas are simple consequences of VP, NPT, DS, and WLF.

Lemma 2 *If player A and player B both play $\langle L, \cdot \rangle$, then $t^1_B \to 0$.*

Lemma 3 *Suppose that player B plays $\langle D, \cdot \rangle$ and then reports b in the second stage, where $0 < b \leq h$ and D is a full-support distribution. If either one of the following conditions holds, then player B is serviced and $t^2_B = 0$:*

1. *$x^1 = 1$*
2. *Player A plays $\langle D', \hat{H} \rangle$, where D' is a full-support distribution.*

We now establish our main lemma: because of the possibility of inconsistent lies, the mechanism cannot charge players with high supertypes. The proof of the main theorem follows from this lemma and details are in the full paper.

Lemma 4 *If player A plays $\langle L, \cdot \rangle$ and B plays $\langle H, \cdot \rangle$, then $t^1_B \to 0$.*

4.4 Incentive Compatibility and Sampling-Based Solutions

As we will see in Section 5, even for simple stochastic welfare maximization problems, there may exist no efficient solutions. To algorithmically implement the desired objective, one may have to approximate the optimal value via sampling, a technique commonly employed in stochastic optimization. In this section, we discuss the impact of such approximation on incentive compatibility.

Theorem 4 *For a given two-stage stochastic optimization problem, suppose that:*

- *there exists a sample average approximation algorithm that finds an ϵ-optimal first-stage decision with prob. $\geq (1 - \xi)$ for any $\epsilon > 0$ and $\xi \in (0, 1)$ in polynomial time;*
- *the exact second-stage optimal decision can be found in polynomial time; and*
- *the worst-case error can be bounded,*

[5] A distribution D is a *full-support* distribution if it has support $(0, h]$ and a cumulative density function that is strictly increasing at every point in its support. Full support distributions play the following role in the proofs: When players report full-support distributions in the first stage, reporting any bid in $(0, h]$ in the second stage is consistent with the behavior of a truthful player.

then an ϵ^*-approximate sequential ex post equilibrium can be algorithmically implemented for any $\epsilon^* \equiv \epsilon^*(\epsilon, \xi) > 0$.

The key idea in the proof is that sampling is required only in the first stage. Hence, to the agents, this additional source of uncertainty only happens in the first stage, when they are interested in maximizing expected utilities. Therefore, the sampling required can be factored into the agents' expected utilities.

Note that the above theorem applies to both multiplicative and additive approximation, with corresponding changes in the incentive guarantees. It demonstrates a trade-off between stronger incentive guarantees and the running time of the sampling-based algorithm. As stochastic welfare maximization involves a mixed-sign objective, in our following result, we focus on additive approximation.

5 A Polynomial Time Implementation for a Class of Stochastic Coverage Cost Problems

To better appreciate the algorithmic challenge posed by stochastic welfare maximization, we now examine the class of stochastic CC (Coverage Cost) problems. This class of problems includes FTM as a special case.

Our approach is based on a combination of the Sample Average Approximation (SAA) method (see, e.g, [12, 17]) and supermodular set function maximization. However, our analysis differs from those found in recent work (e.g. [16, 4]), as we are faced with a mixed-sign objective. To begin, let us first define a single-shot version of the *coverage-cost* problem.

Definition 5 A *coverage cost problem* (CC) consists of three components:

- a set of players $U = \{1, 2, \ldots, n\}$ and a universe of elements E;
- a cost function $c : E \mapsto \mathbb{R}^+$ that assigns a non-negative cost to each element $e \in E$; (we let $c(S) = \sum_{e \in S} c(e)$ for $S \subseteq E$)
- a *service set* $P_s \subseteq E$ for each player s that needs to be constructed in order to serve s, and a value of θ_s of serving s.

The objective is to find $P \subseteq E$ that maximizes *welfare*: $\sum_{s: P_s \subseteq P} \theta_s - c(P)$.

Stochastic CC is defined by extending the CC problem to have two cost functions, c^1 and c^2, for the respective stages. Also, instead of a precise value θ_s of serving agent $s \in U$, in the first stage, only a distribution δ_s is known. The objective is to maximize expected welfare.

By interpreting the universe of elements E as the set of edges in the fixed tree, and the service sets P_s as the (unique) path connecting a node s to the root r of the tree, we see that CC is a generalization of FTM. The following example shows that the generalization is strict:

Example 1 Consider an instance with three players: $U = \{1, 2, 3\}$ and three elements $E = \{a, b, c\}$. Let $P_1 = \{a, b\}$, $P_2 = \{b, c\}$, $P_3 = \{a, c\}$. The cyclic structure entails that this cannot be a FTM instance.

It is formally hard to solve two-stage stochastic CC. The difficulty is not due to a lack of combinatorial structure, as we will show that deterministic CC can indeed be solved in polynomial time. The difficulty is due to the uncertainty in the optimization parameters. We show that it is difficult to solve optimally a stochastic CC instance with only one element (a single-edge FTM instance), even when the distributions are discrete and communicated explicitly as tables of probabilities. The theorem is by reduction from PARTITION ([7]).

Theorem 5 *Maximizing expected welfare for stochastic CC is #P-hard.*

5.1 A Probabilistic Approximation

In view of the above hardness result, we propose a sampling-based solution that *approximates* the expected welfare with high probability. Our algorithm achieves an additive approximation, as multiplicative approximation is unachievable with a polynomial number of samples (due to the mixed-sign objective; example in full paper). The main theorem of this section is as follows:

Theorem 6 *The two-stage stochastic CC problem can be approximated to within an additive error of ϵ in time polynomial in M, $|E|$, and $\frac{1}{\epsilon}$, for all $\epsilon > 0$, where $M = \max_\theta \sum_{s \in U} \theta_s$.*

We now describe the framework for solving two-stage stochastic CC problems. The key structure we will establish and exploit is that optimizations at both stages involve supermodular functions. Recall that a set function $f : 2^N \mapsto \mathbb{R}$ is *supermodular* if for all $S, T \subseteq N$, $f(S) + f(T) \leq f(S \cup T) + f(S \cap T)$. We will show that the expected welfare is supermodular in the set of elements bought in the first stage (see Corollary 1), and that the welfare in the second stage, given the elements bought in the first stage, is also supermodular in the remaining elements. Once these results are established, it is natural to consider the following algorithm:

1. Use the algorithm for supermodular function maximization of [10] to find the optimal first-stage elements to buy. Note that the algorithm needs a value oracle that cannot be implemented in polynomial time. We instead use sampling to approximate the solution value.
2. Given the realized values, use the algorithm of [10] to find the optimal second-stage elements to buy. In this case, the exact value for the value oracle can be found in polynomial time.

In Lemmas 5–8 and Corollary 1, we establish that both the first-stage optimization problem, denoted by $\bar{w}(\cdot)$, and the second-stage optimization problem, denoted by $f_\theta(\cdot)$, involve supermodular functions.

Lemma 5 *For any valuation θ, the function $V_\theta(\cdot)$ defined via:*

$$V_\theta(E') = \sum_{s : P_s \subseteq E'} \theta_s$$

is supermodular in E'.

Lemma 6 *For any valuation θ and any set E_1 of elements bought in the first stage, the second-stage objective $f_\theta(\cdot)$ given by:*

$$f_\theta(E') = V_\theta(E_1 \cup E') - c^2(E')$$

is supermodular in E'.

Lemma 7 *Given any realization θ, the optimal value of the second-stage objective $f^*_\theta(\cdot)$ given by:*

$$f^*_\theta(P) = \max_{F \subseteq E \setminus P} V_\theta(F \cup P) - c^2(F)$$

is supermodular in the set P of elements bought in the first stage.

Lemma 8 *Given any realization θ, the welfare function $w_\theta(\cdot)$ defined via:*

$$w_\theta(P) = f^*_\theta(P) - c^1(P)$$

is supermodular in the set P of elements bought in the first stage.

Corollary 1 *The function $\overline{w}(\cdot) = \mathbb{E}_{\theta \sim \delta}[w_\theta(\cdot)]$ is supermodular.*

Armed with Lemma 8 and Corollary 1, we now address the algorithmic issues of the stochastic CC problem. In the proof of Theorem 5, we have shown that evaluating $\overline{w}(\cdot)$ exactly is #P-hard in general. Fortunately, we can approximate its value in polynomial time, while preserving supermodularity.

Lemma 9 *Let $\mathcal{S}$ be a size $\mathcal{O}(\frac{M^2}{\epsilon^2}|E|)$ set of scenarios drawn from the universe. Let $\hat{w}(\cdot)$ be the sample average approximation of $\overline{w}(\cdot)$ constructed using the samples in $\mathcal{S}$. Then,*

1. *$\hat{w}(\cdot)$ is supermodular;*
2. *for all $F \subseteq E$, $\mathbb{P}\Big[\big|\hat{w}(F) - \overline{w}(F)\big| > \epsilon\Big] \leq o(e^{-|E|})$.*

The proof of Lemma 9 is in the full paper. We now prove Theorem 6.

Proof (Theorem 6). For running time, by using the strongly polynomial-time algorithm of Iwata et al. [10] to perform maximization of supermodular function, the number of function evaluation is bounded by $\mathcal{O}(|E|^5 \log |E|)$. Each function evaluation requires, for each of the $\mathcal{O}((\frac{M^2}{\epsilon^2}|E|))$ samples, finding the second-stage optimal solution given a first-stage solution. The second-stage optimal solution is solved again using supermodular function maximization, and hence each function evaluation takes $\mathcal{O}(|E|^5 \log |E| \times (\frac{M^2}{\epsilon^2}|E|))$ time.

For correctness, by Lemma 9, we can approximate the function $\overline{w}(\cdot)$ by $\hat{w}(\cdot)$ to within an additive error of ϵ' with probability at least $(1 - o(e^{-|E|}))$. As $\hat{w}(\cdot)$ is supermodular by construction, the algorithm of [10] applies.

Given this algorithm, the fact that the second-stage optimization can be solved efficiently using supermodular function maximization, and that the worst-case error of any stochastic CC problem is bounded by $\max\{\max_\theta \sum_{i \in U} \theta_i, c^1(E)\}$, we see that Theorem 4 applies. Thus, the welfare maximizer of stochastic CC problem can be implemented in ϵ-approximate sequential EP equilibrium for any desired $\epsilon > 0$.

Acknowledgments

We thank Kevin Leyton-Brown, David Parkes, and Bob Wilson for fruitful discussion, and especially to Bob for suggesting the term *sequential EP equilibrium*.

References

1. Aaron Archer, Joan Feigenbaum, Arvind Krishnamurthy, Rahul Sami, and Scott Shenker. Approximation and collusion in multicast cost sharing. *Games and Economic Behavior*, 47:36–71, 2004.
2. Dirk Bergemann and Juuso Välimäki. Efficient dynamic auctions. Working paper, 2006.
3. Ruggiero Cavallo, David C. Parkes, and Satinder Singh. Optimal coordinated planning amongst self-interested agents with private state. In *Proc. UAI 2006*, 2006.
4. Moses Charikar, Chandra Chekuri, and Martin Pál. Sampling bounds for stochastic optimization. In *Proc. RANDOM*, pages 257–269, 2005.
5. George B. Dantzig. Linear programming under uncertainty. *Management Science*, 1(3/4):197–206, 1955.
6. Joan Feigenbaum, Christos H. Papadimitriou, and Scott Shenker. Sharing the cost of multicast transmissions. *JCSS*, 63(1):21–41, 2001.
7. Michael R Garey and David S. Johnson. *Computers and Intractability.* W. H. Freeman and Co., New York, 1979.
8. Theodore Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.
9. Nicole Immorlica, David Karger, Maria Minkoff, and Vahab S. Mirrokni. On the costs and benefits of procrastination: approximation algorithms for stochastic combinatorial optimization problems. In *SODA '04*, pages 691–700, 2004.
10. Satoru Iwata, Lisa Fleischer, and Satoru Fujishige. A combinatorial strongly polynomial algorithm for minimizing submodular functions. *JACM*, 48:761–777, 2001.
11. Matthew Jackson. Mechanism theory. In Ulrich Derigs, editor, *Encyclopedia of Life Support Systems.* EOLSS Publishers, Oxford, UK, 2003.
12. Anton J. Kleywegt, Alexander Shapiro, and Tito Homem de Mello. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization*, 12(1):479–502, 2001.
13. Aranyak Mehta, Soctt Shenker, and Vijay Vazirani. Profit-maximizing multicast pricing by approximating fixed points. *J. Algorithms*, 58(2):150–164, 2006.
14. Noam Nisan and Amir Ronen. Algorithmic mechanism design (extended abstract). In *Proc. 31st STOC*, pages 129–140, 1999.
15. Noam Nisan and Amir Ronen. Computationally feasible VCG mechanisms. In *Proc. 1st EC*, pages 242–252, 2000.
16. David B. Shmoys and Chaitanya Swamy. An approximation scheme for stochastic linear programming and its application to stochastic integer programs. *JACM*, 53(6):978–1012, 2006.
17. Chaitanya Swamy and David B. Shmoys. Approximation algorithms for 2-stage stochastic optimization problems. *SIGACT News*, 37(1):33–46, 2006.