Fifty-first Annual Allerton Conference
Allerton House, UIUC, Illinois, USA
October 2 - 3, 2013

A Dynamic VCG Mechanism for Random Allocation Spaces

Maximilian Balandat and Claire J. Tomlin

Abstract— **We construct a dynamic version of the VCG pivot mechanism applicable to private value environments in which the allocation space in each period is random. A key characteristic of this setting is that while the allocations chosen by the coordinator in each stage are based on the actual realization of the random set of feasible allocations, the agent's reports in each step are made prior to this realization. This adds an additional layer to the problem, as players reporting their type now will take into account the distributional characteristics of the feasible allocation sets. We consider a modified social welfare function, in which "the public", regarded as a non-player, also receives utility from allocation decisions. For both the finite and the infinite horizon case, we construct mechanisms that satisfy a suitable notion of incentive compatibility and individual rationality and are weak budget balanced. Finally, we outline how our mechanisms may be applied to the problem of dynamically allocating random goods to a group of players that have private valuations for different item bundles.**

I. Introduction

Mechanisms that implement efficient social choice functions in environments in which participants have private information about their preferences have been studied extensively in the economics literature. A well-known class of such mechanisms are the Vickrey-Clarke-Groves (VCG) mechanisms [1], [2], [3]. VCG mechanisms have been extended to the case of incomplete information in [4], [5], where incentive compatibility is formulated in terms of Bayesian-Nash equilibria. The pivot mechanism [6] is a VCG mechanism in which the players' payoffs equal their respective marginal contributions to the social surplus.

In [7] the idea of the pivot mechanism has been generalized to a dynamic environment with private information. The underlying idea is to design an inter-temporal sequence of transfer payments that allows each player to receive her "flow marginal contribution" in each period. The information asymmetry is with respect to the players' types: at each time, each player privately observes her type. The dynamic pivot mechanism constructed in [7] is efficient, individually rational and ex post incentive compatible. In the related work [8], the authors design an efficient, budget-balanced and Bayesian-Nash incentive compatible dynamic mechanism (which is generally not individually rational). Dynamic mechanism for revenue maximization are investigated in [9].

In this paper, we extend the ideas from [7] in a number of different directions. Firstly, we allow for a publicly observed state that affects payoffs as well as type dynamics. Secondly, we consider a setting in which allocation decisions do not only affect the players, but also create utility for the "public", which does not participate in the mechanism. The mechanism aims to maximize the sum of the player utilities and the "public utility". This can be interpreted as a dynamic version of an affine maximizer problem [10]. Most importantly, we consider random allocation spaces. More precisely, we are interested in problems in which the set of feasible allocations in each period depends on some exogenous random variable, which is realized only after the players have reported their types. The coordinator observes the realization of this random variable (thus the set of feasible allocations), and chooses an allocation such as to maximize social welfare. The players, however, when reporting their types, need to take the distribution of the random variable into account. Finally, we also consider the finite horizon problem, where it is possible to account for player utilities, state dynamics and distributions of feasible allocation sets that are time-varying.

Our main contribution is to construct dynamic direct revelation mechanisms for both the finite and infinite horizon case that satisfy suitable notions of incentive compatibility and individual rationality in a setting where the feasible allocation sets in each period are random and unknown to the players at the time they report their types. This is relevant in engineering problems in which allocation decisions have to be made right after realization of the randomness determining the feasible allocations. One such example is the control of independent storage providers to take up power imbalances in the grid – the dispatch decision has to be made immediately after the power imbalance becomes known.

As an example of a setting in which mechanisms of the type we construct are necessary, we consider the problem of dynamically allocating random items to a group of rational players who have private valuations for different bundles.

Outline: Section II describes our setup. The welfare maximization problem is discussed in section III. Notions of incentive compatibility and individual rationality are formulated in section IV. Our main results are contained in section V; there we propose our mechanisms and analyze their properties. Section VII gives an illustrative example and section VIII concludes the paper.

II. Preliminaries

Consider a discrete-time, private value setting with a finite set $\mathcal{N} = \{1, \dots, N\}$ of players. The stage utility of player $i \in \mathcal{N}$ in period t depends on the player's type $\theta_{i,t} \in \Theta_i$, a publicly observed state $\theta_{0,t} \in \Theta_0$, the publicly

M. Balandat and C. Tomlin are both with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, USA. [balandat,tomlin]@eecs.berkeley.edu

This work has been supported in part by NSF under CPS:ActionWebs (CNS-931843) and CPS:FORCES (CNS-1239166).

observed allocation decision[1] $a_t \in A_t$ and a monetary transfer $p_{i,t} \in \mathbb{R}$. Here A_t is the set of feasible allocation decisions in period t. The type spaces Θ_0 and Θ_i for $i \in \mathcal{N}$ are treated as measurable spaces and assumed to be common knowledge. Define $\Theta := \times_{i=0}^{I} \Theta_i$, $\Theta_{-i} := \times_{j \neq i} \Theta_j$ and $\theta_t := (\theta_{0,t}, \ldots, \theta_{I,t})$. We consider both finite and infinite horizon problems, i.e. $t \in \{0, \ldots, T\}$ with $T \leq \infty$.

The stage utility function $u_{i,t}$ of player i in period t is quasilinear in the monetary transfer and given by

$$u_{i,t}(a_t, p_{i,t}, \theta_{0,t}, \theta_{i,t}) = v_{i,t}(a_t, \theta_{0,t}, \theta_{i,t}) - p_{i,t} \qquad (1)$$

Here $p_{i,t}$ is the payment by player i to the coordinator in period t. If $T < \infty$ each player i receives a terminal utility of the form $g_i(\theta_{0,T}, \theta_{i,T})$ in period T.

Our goal is to construct efficient mechanisms for a slightly more general problem than the classic welfare maximization problem: We assume that, besides providing utility to the players, the coordinator's decisions may also benefit "the public", which is not a participant in the mechanism and has no way of influencing it. For example, in problems of relatively short time scales in power markets, the public could represent the taxpayer, while the players could be market participants. The public's utility is described by a utility function $w_t(a_t, \theta_{0,t})$. If $T = \infty$, we assume the utility functions to be time-invariant, i.e. $v_{i,t}(\cdot,\cdot,\cdot) = v_i(\cdot,\cdot,\cdot)$ and $w_t(\cdot,\cdot) = w(\cdot,\cdot)$ for all $t \geq 0$. All utility functions are assumed to be common knowledge in period $t = 0$.

Let $(\Omega, \mathcal{F}, P)$ be a probability space, and let $(\mathbf{Z}, \mathcal{Z})$ and $(\mathbf{A}, \mathcal{S})$ be measurable spaces. The feasible allocation sets A_t are parametrized by independent[2] random variables $Z_t : \Omega \to \mathbf{Z}$ with distributions μ_{Z_t} (we require the Z_t to be i.i.d. if $T = \infty$). A set-valued function (i.e. a correspondence) $A : \mathbf{Z} \twoheadrightarrow \mathbb{A}$ maps the realization z_t of the random variable Z_t to a set of feasible allocations $A_t = A(z_t) \subset \mathbf{A}$ in $\mathcal{S}$, where $\mathbf{A}$ is the allocation space. In particular, A_t, the set of feasible allocations in period t, is random. This randomness is a major difference between our model and the one in [7], which assumes $A_t = A$ to be fixed and known to every participant in period $t = 0$. The distributions μ_{Z_t} and the correspondence A are assumed common knowledge.

There are many situations in which randomness of the above form may arise. The general class of problems is the allocation of the outcome of some random event across a group of players. In section VII we give an example of a dynamic combinatorial auction, in which randomly realized goods need to be distributed among different players that have private valuations for bundles of items.

Type θ_i of player i is described by a time-homogeneous[3] Markov process on the respective type space Θ_i. For each player $i \in \mathcal{N}$, a common prior $F_i(\theta_{i,0})$ describes the belief about her initial state $\theta_{i,0}$. We work in an independent type setting: A player's type in period t together with the allocation a_t and the state $\theta_{0,t}$ defines a probability distribution for the players's type in period $t+1$, described by a stochastic kernel $F_i(\theta_{i,t+1}; \theta_{0,t}, \theta_{i,t}, a_t)$. Similarly, the evolution of θ_0 is described by a kernel $F_0(\theta_{0,t+1}; \theta_{0,t}, a_t)$. Kernels and common priors are assumed to be independent across players and common knowledge in period $t = 0$.

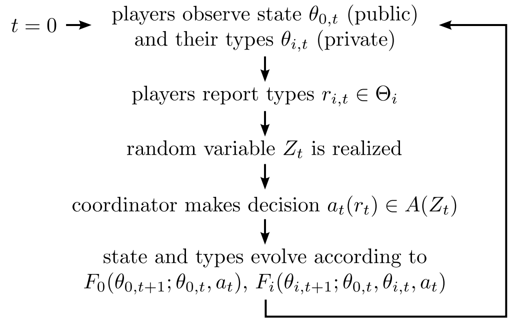

Fig. 1. Information pattern of the mechanism (not including transfers)

At the beginning of each period t, each player i observes her type $\theta_{i,t}$ privately, and all players and the coordinator observe the public state $\theta_{0,t}$. Then each player i reports a type $r_{i,t}$ to the coordinator (where in general $r_{i,t} \neq \theta_{i,t}$, i.e. the players may choose to lie). Then the random variable Z_t is realized and observed publicly. At the end of each period, an allocation $a_t \in A(z_t)$ is chosen by the coordinator and payoffs are realized. Fig. 1 depicts this information pattern.

Remark 1: In our independent types setting, the type $\theta_{i,t+1}$ of player i in period $t+1$ is conditionally independent of $\theta_{j,t}$ for $j \neq i$ given the decision a_t and the state $\theta_{0,t}$. This conditional independence is essential, if it did not hold then the problem would be significantly more difficult. In particular, players could try to affect the evolution of other players' types through that of their own type.

III. Social Welfare Maximization

We are interested in efficient mechanisms that maximize social welfare. If $T = \infty$ we assume all players and the public have a common discount factor $0 < \delta < 1$. The socially efficient policy is obtained by maximizing the expected sum (discounted if $T = \infty$) of the players' utilities and the public's utility over the whole horizon. This can be seen as a dynamic version of an affine maximizer problem [10].

The information structure of the problem is such that the allocation a_t in period t in fact is decided on by the coordinator with respect to the *known* set of feasible allocations $A_t = A(z_t) \subset \mathbf{A}$. For each realization $z \in \mathbf{Z}$, the coordinator's decision in period t is a measurable function $a_t(z) : \Theta \to A(z)$. In order to be able to provide the correct incentives, the coordinator must be able to commit to *allocation policies*:

Definition 1: An *allocation policy* α is a measurable map $\alpha : \mathbf{Z} \times \Theta \to \mathbf{A}$.

An allocation policy α is *admissible* if $\alpha(z, \theta) \in A(z)$ for all $(z, \theta) \in \mathbf{Z} \times \Theta$. Let $\mathcal{A}$ be the set of all admissible policies.

Let $F := \times_{j=0}^{N} F_i$. In order to ensure well-posedness, we assume the following uniform bounds on the expected

[1] The term "allocation" is borrowed from the auction theory literature; more generally a_t is simply the decision the coordinator makes in period t.

[2] This is w.l.o.g. as one could consider inter-temporal correlation of the random variables by suitably augmenting the state space Θ_0.

[3] We could also consider time-inhomogeneous kernels in the finite horizon case, but omit this for brevity.

absolute value of the utilities.

Assumption 1: There exists $K < \infty$ such that for all $t \leq T-2$, $i \in \mathcal{N}$, $\theta \in \Theta$, $z \in \mathbf{Z}$, $a \in A(z)$ and $\alpha \in \mathcal{A}$, $\mathbb{E}_{Z_{t+1}}\left[\int |v_{i,t+1}(\alpha(Z_{t+1},\theta'),\theta_0',\theta_i')|\, dF(\theta';a,\theta)\right] < K$ and $\int |g_i(\theta_0',\theta_i')|\, dF(\theta';a,\theta) < K$. If $T=\infty$ $\exists$ $K_\infty < \infty$ such that $\forall$ $i \in \mathcal{N}$, $\theta \in \Theta$, $z \in \mathbf{Z}$, $a \in A(z)$ and $\alpha \in \mathcal{A}$, $\mathbb{E}_{Z_{t+1}}\left[\int |v_i(\alpha(Z_{t+1},\theta'),\theta_i')|\, dF(\theta';a,\theta)\right] < K$.

Assumption 2: There exists $K_w < \infty$ such that for all $t < T$, $\theta \in \Theta$, $z \in \mathbf{Z}$, $a \in A(z)$ and $\alpha \in \mathcal{A}$, $\mathbb{E}_{Z_{t+1}}\left[\int |w_{t+1}(\alpha(Z_{t+1},\theta'),\theta_0')|\, dF(\theta';a,\theta)\right] < K_w$.

After $Z_t = z_t$ has been realized in period t, the maximal expected social utility to go in the finite horizon case can, under Assumptions 1 and 2, be written as

$$W_t(\theta_t;z_t) = \sup_{a_t\in A(z_t),\{\alpha_s\}_{s=t+1}^{T-1}} \Big\{ J_t(a_t,\theta_t) + \mathbb{E}\big[\textstyle\sum_{s=t+1}^{T-1} J_s(\alpha_s(Z_s,\theta_s),\theta_s) + \sum_{i=1}^N g_i(\theta_{0,T},\theta_{i,T}) \mid \theta_t, a_t\big]\Big\} \tag{2}$$

where $J_t(a_t,\theta_t) = w_t(a_t,\theta_{0,t}) + \sum_{i=1}^N v_{i,t}(a_t,\theta_{0,t},\theta_{i,t})$ is the social stage utility in period t and the expectation is taken over the random variables Z_s for $s > t$ and the evolution of $\theta_{0,t}$ and $\theta_{i,t}$, described by the kernels F_0 and F_i for $i \in \mathcal{N}$. For $T = \infty$, we have under Assumptions 1 and 2 [11]:

$$W_\infty(\theta_t;z_t) = \sup_{a_t\in A(z_t),\{\alpha_s\}_{s=t+1}^{\infty}} \Big\{ J(a_t,\theta_t) + \mathbb{E}\big[\textstyle\sum_{s=t+1}^{\infty} J(\alpha_s(Z_s,\theta_s),\theta_s) \mid \theta_t, a_t\big]\Big\} \tag{3}$$

where $J(a_t,\theta_t) = w(a_t,\theta_{0,t}) + \sum_{i=1}^N v_i(a_t,\theta_{0,t},\theta_{i,t})$.

Our main assumption is the following:

Assumption 3: There exist policies $\boldsymbol{\alpha}_T^* = \{\alpha_t^*\}_{t=0}^{T-1}$ and $\boldsymbol{\alpha}_\infty^* = \{\alpha_t^*\}_{t=0}^{\infty}$ that maximize (2) and (3), respectively.

Assumption 3 gets rid of a number of technical difficulties, in particular potential measurability issues. While these questions are interesting and relevant, we will not consider them in any detail in this paper, and instead focus on designing dynamic mechanisms. In fact, note that Assumption 3 is needed if we are to hope for an efficient mechanism to exist: If the welfare maximization problem under full information does not admit an optimal policy, then we should not expect to be able to construct a mechanism implementing its solution in the incomplete information setting.

Remark 2: One case in which Assumption 3 is satisfied is when Θ and $\mathbf{Z}$ are finite, the allocation spaces $A(z)$ are compact for all $z \in \mathbf{Z}$ and the utility functions are continuous. In general, additional assumptions are required to ensure that optimal policies exist [11].

Under Assumption (3), one may express (2) using dynamic programming recursively as the following Bellman equation:

$$W_t(\theta_t;z_t) = \max_{a_t\in A(z_t)} \big\{ J_t(a_t,\theta_t) + \mathbb{E}\big[W_{t+1}(\theta_{t+1};Z_{t+1}) \,|\, \theta_t, a_t\big]\big\} \tag{4}$$

with terminal condition $W_T(\theta_T;z_T) = \sum_{i=1}^N g_i(\theta_{0,T},\theta_{i,T})$. Similarly, (3) can be written as

$$W_\infty(\theta_t;z_t) = \max_{a_t\in A(z_t)} \big\{ J(a_t,\theta_t) + \delta\,\mathbb{E}\big[W_\infty(\theta_{t+1};Z_{t+1}) \,|\, \theta_t, a_t\big]\big\} \tag{5}$$

IV. Dynamic Mechanisms for Random Allocation Spaces

We restrict our attention to direct revelation mechanisms[4] that implement the socially efficient policies $\boldsymbol{\alpha}_T^*$ and $\boldsymbol{\alpha}_\infty^*$, respectively. A dynamic direct revelation mechanism for random allocation spaces (from now on referred to simply as "mechanism") in every period t asks every player i to report her current state $\theta_{i,t}$. In general, $r_{i,t} \neq \theta_{i,t}$, i.e. the report $r_{i,t}$ need not be truthful. The public history at the beginning of period t is the sequence of observed states and reports, $r_s = (\theta_{0,s}, r_{1,s}, \ldots, r_{N,s})$, realizations of the random variable Z_s, and allocation decisions a_s, and is denoted by $h_t = (r_0, Z_0, a_0, \ldots, r_{t-1}, Z_{t-1}, a_{t-1})$. Similarly, $h_{i,t} = (\theta_{i,0}, r_0, Z_0, a_0, \ldots, \theta_{i,t-1}, r_{t-1}, Z_{t-1}, a_{t-1})$ is player i's private history at the beginning of period t. Denote by $\mathcal{H}_t$ and $\mathcal{H}_{i,t}$ the set of possible public and private (of player i) histories at the beginning of period t, respectively.

Definition 2: A *dynamic direct revelation mechanism for random allocation spaces* is characterized by sequences of allocation policies $\boldsymbol{\alpha} = \{\alpha_t\}_{t=1}^{T-1}$ and monetary transfers $\mathbf{p} = \{p_t\}_{t=1}^{T-1}$, where $\alpha_t \in \mathcal{A}$ and $p_t : \mathcal{H}_t \times \Theta \to \mathbb{R}^N$.

A mechanism as in Definition 2 is *efficient* if $\alpha_t = \alpha_t^*$ for each t, i.e. if the allocation policy is a socially optimal policy as in Assumption 3. A reporting strategy for player i in period t is a mapping from player i's private history to the state space: $r_{i,t} : \mathcal{H}_{i,t} \to \Theta_i$. Note that at the time the transfers p_t are made, the random variable Z_t has already been realized. In particular, the p_i could well be a function of Z_t. However, players can choose their reports based only on their private histories $\mathcal{H}_{i,t}$, which do not include Z_t. As players' incentives should not be based on information unavailable to them, the transfers p_t do not depend on Z_t.

Remark 3: One may be tempted to reformulate the dynamic problem as a completely contingent plan, i.e. by embedding it into a static problem and then invoking standard results from the static mechanism design literature. However, this view does not account for additional strategic possibilities players have in the dynamic model where information arrives over time. In particular, player i in period t bases her report on both her private information and the *past reports of other agents*. As a result, truth-telling in general fails to be a weakly dominant strategy for the static mechanism. Another issue is that the participation constraint must be satisfied not only in the initial period (ex ante), but after each period.

A. Incentive Compatibility

Define $\theta_{-i,t} := (\theta_{0,t}, \ldots, \theta_{i-1,t}, \theta_{i+1,t}, \ldots, \theta_{N,t})$ and $r_{-i,t} := (\theta_{0,t}, r_{1,t}, \ldots, r_{i-1,t}, r_{i+1,t}, \ldots, r_{N,t})$. For a given mechanism, a player using reporting strategy $\mathbf{r}_i = \{r_{i,t}\}_{t=0}^{T-1}$ given strategies $\mathbf{r}_{-i} = \{r_{-i,t}\}_{t=0}^{T-1}$ of the other players expects an overall utility of

$$\mathbb{E}\big[\textstyle\sum_{t=0}^{T-1}\big(v_{i,t}(\alpha^*(Z_t,r_t),\theta_{0,t},\theta_{i,t}) - p_{i,t}(h_t,r_t)\big) + g_i(\theta_{0,T},\theta_{i,T})\big]$$

in the finite horizon case and

$$\mathbb{E}\big[\textstyle\sum_{t=0}^{\infty}\delta^t\big(v_i(\alpha^*(Z_t,r_t),\theta_{0,t},\theta_{i,t}) - p_{i,t}(h_t,r_t)\big)\big]$$

[4] See [12] for a discussion of the revelation principle in dynamic games.

in the infinite horizon case. Note that unlike in (2) and (3), the expectations above are taken also over the random variable Z_t, as players choose their reports prior to the realization of Z_t. For a given efficient mechanism $\{\boldsymbol{\alpha}^*, \mathbf{p}\}$ and reporting strategies $\mathbf{r}_{-i}$, let $V_{i,t}(h_{i,t})$ and $V_{i,\infty}(h_{i,t})$ be the respective finite and infinite horizon value functions of player i, i.e. the expected utilities to go under the optimal reporting strategy:

$$V_{i,t}(h_{i,t}) = \max_{r_{i,t}\in\Theta_i} \mathbb{E}\big[v_{i,t}(\alpha_t^*(z_t, r_{i,t}, r_{-i,t}), \theta_{0,t}, \theta_{i,t}) - p_{i,t}(h_t, (r_{i,t}, r_{-i,t})) + V_{i,t+1}(h_{i,t+1})\big] \quad (6)$$

$$V_{i,\infty}(h_{i,t}) = \max_{r_{i,t}\in\Theta_i} \mathbb{E}\big[v_i(\alpha_t^*(z_t, r_{i,t}, r_{-i,t}), \theta_{0,t}, \theta_{i,t}) - p_i(h_t, (r_{i,t}, r_{-i,t})) + \delta\, V_{i,\infty}(h_{i,t+1})\big] \quad (7)$$

We can now define suitable notions of incentive compatibility for mechanisms of the above form.

Definition 3: An efficient dynamic mechanism for random allocation spaces is *stage-wise ex post incentive compatible* if

$$\begin{aligned}\mathbb{E}_{Z_t}&\big[v_{i,t}(\alpha_t^*(Z_t, \theta_t), \theta_{0,t}, \theta_{i,t}) - p_{i,t}(h_t, \theta_t) + V_{i,t+1}(h_{i,t+1})\big]\\ &\geq \mathbb{E}_{Z_t}\big[v_{i,t}(\alpha_t^*(Z_t, r_{i,t}, \theta_{-i,t}), \theta_{0,t}, \theta_{i,t})\\ &\qquad - p_{i,t}(h_t, (r_{i,t}, \theta_{-i,t})) + V_{i,t+1}(h_{i,t+1})\big] \qquad (8)\end{aligned}$$

for all $i = 1, \ldots, N$, $\theta_t \in \Theta$, $r_{i,t} \in \Theta_i$, $t = 0, \ldots, T-1$ and histories $h_{i,t} \in \mathcal{H}_{i,t}$.

Definition 4: An efficient dynamic mechanism for random allocation spaces is *periodic ex post incentive compatible* if

$$\begin{aligned}\mathbb{E}_{Z_t}&\big[v_i(\alpha_t^*(Z_t, \theta_t), \theta_{0,t}, \theta_{i,t}) - p_i(h_t, \theta_t) + \delta\, V_{i,\infty}(h_{i,t+1})\big]\\ &\geq \mathbb{E}_{Z_t}\big[v_i(\alpha_t^*(Z_t, r_{i,t}, \theta_{-i,t}), \theta_{0,t}, \theta_{i,t}) - p_i(h_t, (r_{i,t}, \theta_{-i,t}))\\ &\qquad + \delta\, V_{i,\infty}(h_{i,t+1})\big] \qquad (9)\end{aligned}$$

for all i, $\theta_t \in \Theta$, $r_{i,t} \in \Theta_i$, $t \geq 0$ and histories $h_{i,t} \in \mathcal{H}_{i,t}$. Definitions 3 and 4 characterize a mechanism as incentive compatible if truth-telling is a best response in expectation (w.r.t the random variables Z_t), regardless of the history and the current state of the other players, provided that all other players report truthfully. Note that without the expectation over Z_t, Definition 4 would be the notion of ex post incentive compatibility from [7]. The qualifications "stage-wise" and "periodic" mean that incentive compatibility is ex post w.r.t. to all signals received until period t (*not* including Z_t), but not w.r.t to signals received in later periods. This is because the additional information revealed in future periods will generally render the expectation-maximizing decision suboptimal in hindsight.

Remark 4: Observe that a stage-wise / periodic ex post incentive compatible mechanism guarantees that a reporting strategy profile of the agents is an equilibrium for all possible beliefs about the *agent types*. The assumption that the distributions μ_{Z_t} and the kernels F_i are common knowledge on the other hand is crucial.

B. Individual Rationality

In a similar way, we can define suitable notions of individual rationality. For this, we will need to consider the social utilities in the absence of player i:

$$\begin{aligned}W_{-i,t}(\theta_t; z_t) = &\sup_{a_t\in A(z_t), \{\alpha_s\}_{s=t+1}^{T-1}} \Big\{J_{-i,t}(a_t, \theta_t)\\ &+ \mathbb{E}\big[\textstyle\sum_{s=t+1}^{T-1} J_{-i,s}(\alpha_s(Z_s, \theta_s), \theta_s) + \sum_{j\neq i} g_j(\theta_{0,T}, \theta_{j,T}) \,|\, \theta_t, a_t\big]\Big\}\end{aligned} \quad (10)$$

with $J_{-i,t}(a_t, \theta_t) = w_t(a_t, \theta_{0,t}) + \sum_{j\neq i}^N v_{j,t}(a_t, \theta_{0,t}, \theta_{j,t})$ and

$$\begin{aligned}W_{-i,\infty}(\theta_t; z_t) = &\sup_{a_t\in A(z_t), \{\alpha_s\}_{s=t+1}^{\infty}} \Big\{J_{-i}(a_t, \theta_t)\\ &+ \mathbb{E}\big[\textstyle\sum_{s=t+1}^{\infty} J_{-i}(\alpha_s(Z_s, \theta_s), \theta_s) \mid \theta_t, a_t\big]\Big\}\end{aligned} \quad (11)$$

with $J_{-i}(a_t, \theta_t) = w(a_t, \theta_{0,t}) + \sum_{j\neq i}^N v_j(a_t, \theta_{0,t}, \theta_{j,t})$.

Denote the socially optimal policies when player i is excluded from the mechanism by $\boldsymbol{\alpha}^*_{-i,T} = \{\alpha^*_{-i,t}\}_{t=0}^{T-1}$ and $\boldsymbol{\alpha}^*_{-i,\infty} = \{\alpha^*_{-i,t}\}_{t=0}^{\infty}$, respectively.

Suppose now that after each history h_t, each player may permanently opt out from the mechanism. Let $O_{i,t}(h_{i,t})$ ($O_{i,\infty}(h_{i,t})$) be the expected overall utility player i receives if she opts out, i.e. if the socially efficient policy $\boldsymbol{\alpha}^*_{-i,T}$ ($\boldsymbol{\alpha}^*_{-i,\infty}$) for the remaining players is implemented. For a player i to have an incentive to continue participating in the mechanism, her expected remaining overall utility must be at least as high as $O_{i,t}(h_{i,t})$ ($O_{i,\infty}(h_{i,t})$). This motivates the following definitions:

Definition 5: A dynamic mechanism for random allocation spaces is *stage-wise ex post individually rational* if $V_{i,t}(h_{i,t}) \geq O_{i,t}(h_{i,t})$ for all $h_{i,t} \in \mathcal{H}_{i,t}$ and all $0 \leq t < T$.

Definition 6: A dynamic mechanism for random allocation spaces is *periodic ex post individually rational* if $V_{i,\infty}(h_{i,t}) \geq O_{i,\infty}(h_{i,t})$ for all $h_{i,t} \in \mathcal{H}_{i,t}$ and all $t \geq 0$.

V. The Mechanism

In this section, we construct efficient dynamic direct revelation mechanisms for feasible allocation spaces, for both the finite horizon and infinite horizon problem. It is well known that, even in the static case with quasi-linear preferences, it is impossible to construct a Bayesian-Nash incentive compatible mechanism that achieves efficiency, (strong) budget balance and (interim) individual rationality [13]. As in [7], we will construct mechanisms that, besides efficiency and incentive compatibility (cf. Definitions 3 and 4), ensure individual rationality (cf. Definitions 5 and 6), but which are in general not (strongly) budget balanced.

A. Marginal Contributions

The (expected) *marginal contribution* $M_{i,t}(\theta_t)$ of player i in period t is the expected utility she contributes to society:

$$M_{i,t}(\theta_t) := \mathbb{E}_{Z_t}[W_t(\theta_t; Z_t) - W_{-i,t}(\theta_t; Z_t)] \quad (12)$$

If $T = \infty$ this is a stationary quantity:

$$M_{i,\infty}(\theta_t) := \mathbb{E}_{Z_t}[W_\infty(\theta_t; Z_t) - W_{-i,\infty}(\theta_t; Z_t)] \quad (13)$$

One of the key observations of [7] is that if a player can secure her marginal contribution in every continuation game of the mechanism, then she should be able to incur her *flow marginal contribution* $m_{i,t}(\theta_t)$ ($m_{i,\infty}(\theta_t)$) in every period.

The respective flow marginal contributions in the setting of random feasible allocation sets are defined by

$$m_{i,t}(\theta_t) = M_{i,t}(\theta_t) - \mathbb{E}\left[M_{i,t+1}(\theta_{t+1}) \,|\, \theta_t\right] \quad (14)$$

$$m_{i,\infty}(\theta_t) = M_{i,\infty}(\theta_t) - \delta\, \mathbb{E}\left[M_{i,\infty}(\theta_{t+1}) \,|\, \theta_t\right] \quad (15)$$

To alleviate notational burden, define $\aleph_t^* = \alpha_t^*(Z_t, \theta_t)$ and $\aleph_{-i,t}^* = \alpha_{-i,t}^*(Z_t, \theta_{-i,t})$. Using (12) and (13), respectively, the flow marginal contributions can be expressed as

$$\begin{aligned} m_{i,t}(\theta_t) = \mathbb{E}\Big[& J_t(\aleph_t^*, \theta_t) - J_{-i,t}(\aleph_{-i,t}^*, \theta_t) \\ & + \big(\mathbb{E}\big[W_{-i,t+1}(\theta_{t+1}, z_{t+1}) \,|\, \aleph_t^*, \theta_t\big] \\ & - \mathbb{E}\big[W_{-i,t+1}(\theta_{t+1}, z_{t+1}) \,|\, \aleph_{-i,t}^*, \theta_t\big]\big) \,|\, \theta_t\Big] \end{aligned} \quad (16)$$

$$\begin{aligned} m_{i,\infty}(\theta_t) = \mathbb{E}\Big[& J(\aleph_t^*, \theta_t) - J_{-i}(\aleph_{-i,t}, \theta_t) \\ & + \delta\big(\mathbb{E}\big[W_{-i,\infty}(\theta_{t+1}, z_{t+1}) \,|\, \aleph_t^*, \theta_t\big] \\ & - \mathbb{E}\big[W_{-i,\infty}(\theta_{t+1}, z_{t+1}) \,|\, \aleph_{-i,t}^*, \theta_t\big]\big) \,|\, \theta_t\Big] \end{aligned} \quad (17)$$

B. Monetary Transfers

The basic idea behind the Clarke pivot rule is to make the players internalize the externalities they impose on society in each period. This can be achieved by designing, for each player i, a monetary transfer policy $p_{i,t} : \Theta \to \mathbb{R}$ ($p_{i,\infty} : \Theta \to \mathbb{R}$) such that the resulting expected flow net utility matches the flow marginal contribution:

$$p_{i,t}^*(\theta_t) := \mathbb{E}\left[v_{i,t}(\alpha_t^*(Z_t, \theta_t), \theta_{0,t}, \theta_{i,t})\right] - m_{i,t}(\theta_t) \quad (18)$$

$$p_{i,\infty}^*(\theta_t) := \mathbb{E}\left[v_i(\alpha_t^*(Z_t, \theta_t), \theta_{0,t}, \theta_{i,t})\right] - m_{i,\infty}(\theta_t) \quad (19)$$

Here the expectations are over the random variable Z_t. Observe that the transfers $p_{i,t}^*(\theta_t)$ and $p_{i,\infty}^*(\theta_t)$ depend only on the report θ_t and not on the entire public history h_t. Let $\mathcal{M} = \{\boldsymbol{\alpha}^*, \mathbf{p}^*\}$ and $\mathcal{M}_\infty = \{\boldsymbol{\alpha}_\infty^*, \mathbf{p}_\infty^*\}$ denote the finite and infinite horizon mechanism, respectively. Our first assertion is that both mechanisms $\mathcal{M}$ and $\mathcal{M}_\infty$ need not be subsidized.

Lemma 1: The mechanisms $\mathcal{M}$ and $\mathcal{M}_\infty$ are weak budget-balanced, i.e. $\sum_{i\in\mathcal{N}} p_{i,t}(\theta_t) \geq 0$ for all t and $\theta_t \in \Theta$.

Proof: We prove this for $\mathcal{M}$, the argument for $\mathcal{M}_\infty$ is almost identical. For $\alpha_t', \alpha_{-i,t}' \in \mathcal{A}$, let $\aleph_t' = \alpha_t'(Z_t, \theta_t)$ and $\aleph_{-i,t}' = \alpha_{-i,t}'(Z_t, \theta_{-i,t})$. By definition, $\alpha_{-i,t}^*$ is an optimal policy maximizing the expected social utility of all players excluding player i. That is,

$$\begin{aligned} & \mathbb{E}\Big[J_{-i,t}(\aleph_{-i,t}^*, \theta_t) + \mathbb{E}\big[W_{-i,t+1}(\theta_{t+1}, Z_{t+1}) \,|\, \aleph_{-i,t}^*, \theta_t\big] \,|\, \theta_t\Big] \\ & \geq \mathbb{E}\Big[J_{-i,t}(\aleph_t', \theta_t) + \mathbb{E}\big[W_{-i,t+1}(\theta_{t+1}, Z_{t+1}) \,|\, \aleph_t', \theta_t\big] \,|\, \theta_t\Big] \end{aligned} \quad (20)$$

for all $\alpha_t' \in \mathcal{A}$, $\theta_t \in \Theta$ and $t \in \{0, \ldots, T-1\}$. Using (16) one can express the transfers $p_{i,t}^*$ in terms of the stage utilities and the social continuation values:

$$\begin{aligned} p_{i,t}^*(\theta_t) = \mathbb{E}\Big[& J_{-i,t}(\aleph_{-i,t}^*, \theta_t) - J_{-i,t}(\aleph_t^*, \theta_t) \\ & + \big(\mathbb{E}\big[W_{-i,t+1}(\theta_{t+1}, Z_{t+1}) \,\big|\, \aleph_{-i,t}^*, \theta_t\big] \\ & - \mathbb{E}\big[W_{-i,t+1}(\theta_{t+1}, Z_{t+1}) \,\big|\, \aleph_t^*, \theta_t\big]\big) \,\big|\, \theta_t\Big] \end{aligned} \quad (21)$$

In particular, for $\alpha_t' = \alpha_t^*$ we conclude from (20) that $p_{i,t}^*(\theta_t) \geq 0$ for all $\theta_t \in \Theta$, $i \in \mathcal{N}$ and $t < T$. ■

C. Incentive Compatibility

Our main result is the following:

Theorem 1: The mechanisms $\mathcal{M}$ and $\mathcal{M}_\infty$ are stage-wise and periodic ex post incentive compatible, respectively.

Proof: Again we prove only the finite horizon case $\mathcal{M}$, as the argument for $\mathcal{M}_\infty$ is essentially the same. We are to show that truth-telling is incentive compatible in expectation (w.r.t. Z_t) for every player $i \in \mathcal{N}$ in every period $t \in \{0, \ldots, T-1\}$, provided that all other players report truthfully. By the principle of optimality, it suffices to show that each player i in each period t incurs her marginal contribution as her reward to go if she reports truthfully, provided that all other players do. That is, we are to show that, for all $r_{i,t} \in \Theta_i$, $\theta_{-i,t} \in \Theta_{-i}$ and $t \in \{0, \ldots, T-1\}$,

$$\begin{aligned} & \mathbb{E}\Big[v_{i,t}(\alpha_t^*(Z_t, \theta_t), \theta_{0,t}, \theta_{i,t}) - p_{i,t}^*(\theta_t) \\ & \quad + \mathbb{E}\big[M_{i,t+1}(\theta_{t+1}) \,|\, \alpha_t^*(Z_t, \theta_t), \theta_t\big]\Big] \\ & \geq \mathbb{E}\Big[v_{i,t}(\alpha_t^*(Z_t, r_{i,t}, \theta_{-i,t}), \theta_{0,t}, \theta_{i,t}) - p_{i,t}^*(r_{i,t}, \theta_{-i,t}) \\ & \quad + \mathbb{E}\big[M_{i,t+1}(\theta_{t+1}) \,|\, \alpha_t^*(Z_t, r_{i,t}, \theta_{-i,t}), \theta_t\big]\Big] \end{aligned} \quad (22)$$

Let $t \in \{0, \ldots, T-1\}$, $r_{i,t} \in \Theta_i$ and $\theta_{-i,t} \in \Theta_{-i}$ be arbitrary and denote by $L(\theta_t, t)$ and $R(r_{i,t}, \theta_{-i,t}, t)$ left and right hand side of (22), respectively. The transfer $p_{i,t}^*$ was constructed in (21) exactly so that $L(\theta_t, t) = M_{i,t}(\theta_t) = \mathbb{E}\left[W_t(\theta_t; Z_t) - W_{-i,t}(\theta_t; Z_t)\right]$. Substituting $p_{i,t}^*$ from (21) we can write $R(r_{i,t}, \theta_{-i,t}, t)$ after rearranging terms as

$$\begin{aligned} \mathbb{E}\Big[& J_t(\alpha_t^*(Z_t, r_{i,t}, \theta_{-i,t}), \theta_t) - J_{-i,t}(\alpha_t^*(Z_t, r_{i,t}, \theta_{-i,t}), \theta_t) \\ & + \big(\mathbb{E}\big[W_{t+1}(\theta_{t+1}, Z_{t+1}) \,|\, \alpha_t^*(Z_t, r_{i,t}, \theta_{-i,t}), \theta_t\big] \\ & - \mathbb{E}\big[W_{-i,t+1}(\theta_{t+1}, Z_{t+1}) \,|\, \alpha_{-i,t}^*(Z_t, \theta_{-i,t}), \theta_t\big]\big) \,\big|\, \theta_t\Big] \end{aligned}$$

Using the definition of $\alpha_{-i,t}^*$ and $W_{-i,t}$ we get

$$\begin{aligned} R(r_{i,t}, \theta_{-i,t}, t) = \mathbb{E}\Big[& J_t(\alpha_t^*(Z_t, r_{i,t}, \theta_t), \theta_t) - W_{-i,t}(\theta_t; Z_t) \\ & + \mathbb{E}\big[W_{t+1}(\theta_{t+1}, Z_{t+1}) \,|\, \alpha_t^*(Z_t, r_{i,t}, \theta_{-i,t}), \theta_t\big]\Big] \end{aligned} \quad (23)$$

From (23) and the definitions of α_t^* and W_t we conclude that $L(\theta_t, t) \geq R(r_{i,t}, \theta_{-i,t}, t)$, i.e. that (22) holds. ■

D. Individual Rationality

Let $A^{(N)} : \mathbf{Z} \twoheadrightarrow \mathbb{A}$ denote the correspondence parametrizing the set of feasible allocations when N players participate in the mechanism. In order to be able to show individual rationality, we need two additional Assumptions.

Assumption 4: If $N, M \in \mathbb{N}$ with $N \leq M$, then $A^{(N)}(z) \subseteq A^{(M)}(z)$ for all $z \in \mathbf{Z}$.

Assumption 5: $\mathbb{E}[v_{i,t}(\alpha_{-i,t}^*(Z_t, \theta_{-i,t}), \theta_{0,t}, \theta_{i,t})] \geq 0$ for all $\theta_t \in \Theta$ and $\mathbb{E}[g_i(\theta_{0,T}, \theta_{i,t}) \,|\, \alpha_{-i,t}^*, \theta_{T-1}] \geq 0$ for all $\theta_{T-1} \in \Theta$, for all $i \in \mathcal{N}$ and $t < T < \infty$. If $T = \infty$, $\mathbb{E}[v_i(\alpha_{-i,\infty}^*(Z_t, \theta_{-i,t}), \theta_{0,t}, \theta_{i,t})] \geq 0$, $\forall\, \theta_t \in \Theta$ and $i \in \mathcal{N}$.

Assumption 4 is a type of "choice set monotonicity". Intuitively, it says that removing agents from the mechanism will never enlarge the coordinator's decision set. Assumption 5 ensures that there are no negative externalities.

Theorem 2: Suppose Assumptions 4 and 5 hold. Then the mechanisms $\mathcal{M}$ and $\mathcal{M}_\infty$ are stage-wise and periodic ex post individually rational, respectively.

Proof: Again we prove only the case $T < \infty$. We are to show that $V_{i,t}(h_{i,t}) \geq O_{i,t}(h_{i,t})$ for all $h_{i,t} \in \mathcal{H}_{i,t}$. Recall from the proof of Theorem 2 that (by incentive compatibility) $V_{i,t}(h_{i,t}) = \mathbb{E}\left[W_t(\theta_t; Z_t) - W_{-i,t}(\theta_t; Z_t)\right]$. Therefore

$$\begin{aligned}
&V_{i,t}(h_{i,t}) - O_{i,t}(h_{i,t}) \\
&= \mathbb{E}\big[\textstyle\sum_{s=t+1}^{T-1} J_s(\aleph_s^*, \theta_s) + \sum_{i=1}^N g_i(\theta_{0,T}, \theta_{i,T}) \,|\, \theta_t\big] \\
&\quad - \mathbb{E}\big[\textstyle\sum_{s=t+1}^{T-1} J_{-i,s}(\aleph_{-i,s}^*, \theta_{-i,s}) + \sum_{j\neq i} g_j(\theta_{0,T}, \theta_{i,T}) \,|\, \theta_t\big] \\
&\quad - \mathbb{E}\big[\textstyle\sum_{s=t+1}^{T-1} v_{i,s}(\aleph_{-i,s}^*, \theta_{0,s}, \theta_{i,s}) + g_i(\theta_{0,T}, \theta_{i,T}) \,|\, \theta_t\big] \\
&= \mathbb{E}\big[\textstyle\sum_{s=t+1}^{T-1} J_s(\aleph_s^*, \theta_s) + \sum_{i=1}^N g_i(\theta_{0,T}, \theta_{i,T}) \,|\, \theta_t\big] \\
&\quad - \mathbb{E}\big[\textstyle\sum_{s=t+1}^{T-1} J_s(\aleph_{-i,s}^*, \theta_{-s}) + \sum_{i=1}^N g_i(\theta_{0,T}, \theta_{i,T}) \,|\, \theta_t\big]
\end{aligned}$$

which is nonnegative by Assumptions 4 and 5. ■

Note that Assumptions 4 and 5 are sufficient conditions, i.e. individual rationality may also hold if they are not satisfied.

VI. Implementation Challenges

A. Computational Complexity

In order to implement the mechanism from the previous sections, the coordinator needs to be able to do the following:

1) Given a report r_t and realization $Z_t = z_t$, determine the socially optimal allocation $\alpha_t^*(z_t, r_t)$ or $\alpha_\infty^*(z_t, r_t)$
2) Compute the transfers $p_{i,t}^*$ or $p_{i,\infty}^*$ as a function of the players' reports r_t

In the finite horizon case, the above tasks amount to solving the stochastic optimal control problems (2) and (10) for each $i \in \mathcal{N}$ and $0 \leq t < T-1$. In principle, this can be done by solving the dynamic programming recursion (4) with terminal value functions $W_T(\cdot\,;\cdot)$ and $W_{-i,T}(\cdot\,;\cdot)$ given by $W_T(\theta_T; z_T) = \sum_{i=1}^N g_i(\theta_{0,T}, \theta_{i,T})$ and $W_{-i,T}(\theta_T; z_T) = \sum_{j\neq i} g_j(\theta_{0,T}, \theta_{j,T})$, respectively. Note though that in order to compute the expectation in (21), the coordinator not only needs to know the optimal allocation $a^*(r_t) = \arg\max_{a\in A(z_t)} J_t(a_t, r_t) + \mathbb{E}[W_{t+1}(r_{t+1}; z_{t+1}) \,|\, r_t, a_t]$ but the full optimal *policy* α_t^*.

Depending on the parameters of the problem, i.e. the players' utility functions, the priors $F_i(\theta_{i,0})$, $F_{Z_t}(z_t)$, the stochastic kernels $F_i(\theta_{i,t+1}; \theta_{0,t}, \theta_{i,t}, a_t)$ and the characteristics of the feasible allocation sets (described by the correspondence A), solving these stochastic optimal control problems is generally a very hard task. If Θ and $\mathbf{Z}$ are finite, then the problem is a finite state Markov Decision Process and thus tractable if the cardinality of $\Theta \times \mathbf{Z}$ is sufficiently small. In general, however, there is little one can say about computational tractability given our very general model. These questions should be investigated with respect to the particular application at hand, and is left for future work. Finally, let us emphasize that computational complexity is not a problem exclusive to our model, but one of the fundamental problems in mechanism design [14].

B. Informational Requirements

Aside from purely computational considerations, another potential problem with our mechanism is that it has very high informational requirements. Specifically, among other things, we assume that the players' cost functions, the distributions μ_{Z_t} of the random variables Z_t as well as the priors $F_i(\theta_{i,0})$ and the dynamics $F_i(\theta_{i,t+1}; \theta_{0,t}, \theta_{i,t}, a_t)$ are common knowledge. In many applications, these assumptions are too restrictive. Again, this problem is not exclusive to our approach, but characteristic of many optimal mechanisms. Another issue with optimal mechanisms is that they are often not "robust". That is, they can be very sensitive to details of the environment that will often not be fully known in practice. "Robust Mechanism Design" [15] tries to relax the informational requirements and design simpler mechanisms that trade of optimality for robustness. The literature on designing *robust dynamic* mechanisms is still small, one contribution is the recent paper [16].

VII. Example: Dynamic Combinatorial Allocation of Random Goods

In this section, we show how our mechanism may be employed to implement a socially efficient allocation in a problem of dynamically allocating random goods to a group of players that have private valuations for bundles of items. This can be seen as a combinatorial auction [17] in a dynamic setting in which items arrive randomly over time[5].

A. A Simple Model

Let $\mathbf{Z}$ be a finite set of items ("goods") with $\emptyset \in \mathbf{Z}$, and let μ_Z be a probability distribution over $\mathbf{Z}$. In each period t, a random item $Z_t \in \mathbf{Z}$ is realized according to the distribution μ_Z. For example, let $\mathbf{Z} = \{\emptyset, B, Y, H\}$ with $B =$ Barley, $Y =$ Yeast and $H =$ Hops.

Let $\mathcal{N}$ be a set of N players ("producers"), each of which has private valuations of all[6] item bundles (i.e. all subsets of $\mathbf{Z}$). That is, player i's type $\theta_{i,t}$ in period t is a vector of non-negative valuations, one for each bundle $\mathbf{b} \in 2^{\mathbf{Z}}$, so $\Theta_i = \mathbb{R}^{|2^{\mathbf{Z}}|}$ for all $i \in \mathcal{N}$. We assume that $\theta_{i,t}(\emptyset) = 0$ for all players i in all periods t. In our example, the type $\theta_{i,t}(\mathbf{b})$ could represent the profit producer i can make in period t when using the items in $\mathbf{b}$ as inputs. The items can complement each other; in our example it is clear that $\theta_{i,t}(\{B,Y,H\}) \gg \theta_{i,t}(\{B\}) + \theta_{i,t}(\{Y\}) + \theta_{i,t}(\{H\})$.

Suppose that the valuations $\theta_{i,t}$ evolve over time, with dynamics described by stochastic kernels $F_i(\theta_{i,t+1}; \theta_{i,t})$. A deterministic part of F_i in our example describes discounting over time, while the stochastic part could represent the effect of the current weather on the quality of B, Y and H.

We assume that which items have previously been allocated to the different players is public information. To this end, let $\theta_{0,t} \in \times_{i=1}^N 2^{\mathbf{Z}}$ be the publicly observable state, so that $\theta_{0,t}(i) \subset \mathbf{Z}$ is the set of items that have been allocated

[5]Note that this is different from *dynamic implementations* of combinatorial auctions – our setting is inherently dynamic.

[6]We could restrict our attention to a subclass of bundles in order to reduce the dimension of the problem, but refrain from it here for simplicity.

to player i prior to period t. Conditioned on a realization $z_t \in \mathbf{Z}$ of Z_t, the one step dynamics of $\theta_{0,t}$ are deterministic and given by $\theta_{0,t+1}(i) = \theta_{0,t}(i) \cup a_{i,t}$, where $a_t \in A_t = A(z_t) = \{\{z_t, \emptyset, \dots\}, \{\emptyset, z_t, \emptyset, \dots,\}, \dots, \{\emptyset, \dots, \emptyset, z_t\}\}$ is the allocation decision made by the mechanism in period t. Note that each player can use at most one item of the same type, and that at most a single item can be allocated in each period, so $a_{i,t} \neq \emptyset$ for at most one $i \in \mathcal{N}$.

Consider a finite horizon setting with $T < \infty$ and $g_i = v_i$ for all $i \in \mathcal{N}$. In each period t, players receive a utility $v_i(\theta_{0,t}, \theta_{i,t}) = \theta_{i,t}(\theta_{0,t}(i))$ based on the items allocated to them in previous periods. While the allocation $a_{i,t}$ does not affect the stage utility v_i in period t, it does affect the utilities in subsequent periods through its effect on $\theta_{0,s}(i)$ for $s > t$.

The reason why implementing the socially efficient allocation for the above problem cannot be done naïvely and requires the design of a mechanism is the same as in the static auction case: players in general will have an incentive to misreport their valuations $\theta_{i,t}$ for the different item bundles in order to maximize their expected overall utility. A static mechanism is not applicable because there is inter-temporal coupling through the allocation decisions. These are based on the players' reported valuations, which in turn depend on private information that arrives over time (as described by the type dynamics F_i). The mechanism from [7] is also not applicable, since it assumes the set of feasible allocations $A_t = A$ to be fixed and known a priori.

B. Implementing the Mechanism

The Bellman equation (4) for the social welfare maximization for the above model reads

$$\begin{aligned} W_t(\theta_t; z_t) = \max_{a_t \in A(z_t)} \Big\{ &\textstyle\sum_{i=1}^N v_i(\theta_{0,t}, \theta_{i,t}) \\ &+ \textstyle\sum_{z \in \mathbf{Z}} \mu_Z(z)\, \mathbb{E}_{\theta_{i,t+1}} \big[W_{t+1}(\theta_{t+1}; Z_{t+1}) \,|\, \theta_t, a_t \big] \Big\} \end{aligned} \tag{24}$$

where $W_T(\theta_T; z_T) = \sum_{i=1}^N v_i(\theta_{0,T}, \theta_{i,T})$. Since both $\mathbf{Z}$ and $A(z_t)$ are finite, the main challenge in (24) is the computation of the expectation over the private types $\theta_{i,t+1}$ (note that a_t only affects $\theta_{0,t+1}$ in a deterministic way, but not $\theta_{i,t+1}$).

Having computed W_∞, for each $z_t \in \mathbf{Z}$ the optimal allocation $\alpha_t^*(z_t, \theta_t)$ as a function of θ_t is given as the maximizer in (24). The value functions $W_{-i,\infty}$ defined in (10) and the associated policies $\alpha_{-i,t}^*(z_t, \theta_t)$ can be determined analogously and then be used to compute the transfers

$$\begin{aligned} p_{i,t}^*(\theta_t) = \textstyle\sum_{z \in \mathbf{Z}} \mu_Z(z) \Big\{ & \\ &\textstyle\sum_{j \neq i} \big(v_i(\alpha_{-i,t}^*(\theta_t, z), \theta_t) - v_i(\alpha_t^*(\theta_t, z), \theta_t) \big) \\ &+ \textstyle\sum_{z' \in \mathbb{Z}} \mu_Z(z') \Big[\mathbb{E}_{\theta_{t+1}} \big[W_{-i,t+1}(\theta_{t+1}, z') \,|\, \alpha_{-i,t}^*(\theta_t, z), \theta_t \big] \\ &\qquad - \mathbb{E}_{\theta_{t+1}} \big[W_{-i,t+1}(\theta_{t+1}, z') \,|\, \alpha_{i,t}^*(\theta_t, z), \theta_t \big] \Big] \Big\} \end{aligned}$$

With the optimal policies $\{\alpha_t^*\}_{t=0}^{T-1}$ and transfers $\{p_t^*\}_{t=0}^{T-1}$ at hand, one can implement a dynamic direct revelation mechanism as described in section IV. By design, the resulting mechanism is stage-wise ex post incentive compatible. Recalling that $v_i \geq 0$, it is easy to verify that Assumptions 4 and 5 are satisfied, which means that the mechanism for this problem is also stage-wise ex post individually rational.

VIII. Conclusion and Future Work

In this paper we considered the problem of maximizing the social welfare of a group of rational players with private information in a dynamic setting with private values, in which the feasible allocation sets are random. Our notion of social welfare included utility for "the public" created by allocation decisions. We have constructed efficient dynamic direct revelation mechanisms for both the finite and infinite horizon problem, which satisfy suitable notions of incentive-compatibility and, under additional assumptions, individual rationality. Moreover, we have shown that the mechanisms are weak budget balanced. Finally, we illustrated how our mechanisms may be applied to the problem of dynamically allocating random goods to a group of players that have private valuations for bundles of items.

A common theme in the design of optimal mechanisms, the main issues that limit the use of our proposed mechanisms in practice are their high computational complexity and informational requirements. Going forward, we plan to employ techniques from robust (dynamic) mechanism design to construct simpler mechanisms that are easier to implement, pose less restrictive assumptions on information available to agents and coordinator, and are less sensitive to changes of or uncertainties in the environment.

References

[1] T. Groves, "Incentives in teams," *Econometrica*, vol. 41, no. 4, pp. 617–631, 07 1973.

[2] E. H. Clarke, "Multipart pricing of public goods," *Public Choice*, vol. 11, pp. 17–33, 10 1971.

[3] W. Vickrey, "Coutnerspeculation, auctions, and competitive sealed tenders," *The Journal of Finance*, vol. 16, no. 1, pp. 8–37, 1961.

[4] C. d'Aspremont and L.-A. Gérard-Varet, "Incentives and incomplete information," *Journal of Public Economics*, vol. 11, no. 1, pp. 25–45, 2 1979.

[5] K. J. Arrow, *The Property Rights Doctrine and Demand Revelation Under Incomplete Information*, ser. Economics and human welfare. Academic Press, 1979.

[6] J. Green and J.-J. Laffont, "Characterization of satisfactory mechanisms for the revelation of preferences for public goods," *Econometrica*, vol. 45, no. 2, pp. 427–438, 03 1977.

[7] D. Bergemann and J. Välimäki, "The dynamic pivot mechanism," *Econometrica*, vol. 78, no. 2, pp. 771–789, 2010.

[8] S. Athey and I. Segal, "An efficient dynamic mechanism," working paper, Sep 2012.

[9] A. Pavan, I. Segal, and J. Toikka, "Dynamic mechanism design: A myersonian approach," *Econometrica (accepted)*, 2013.

[10] N. Nisan and A. Ronen, "Computationally feasible vcg mechanisms," *J. Artif. Int. Res.*, vol. 29, no. 1, pp. 19–47, May 2007.

[11] N. Stokey, R. Lucas, and E. Prescott, *Recursive Methods in Economic Dynamics*. Harvard University Press, 1989.

[12] R. B. Myerson, "Multistage games with communication," *Econometrica*, vol. 54, no. 2, pp. 323–358, 03 1986.

[13] R. B. Myerson and M. A. Satterthwaite, "Efficient mechanisms for bilateral trading," *Journal of Economic Theory*, vol. 29, no. 2, pp. 265 – 281, 1983.

[14] V. Conitzer and T. Sandholm, "Complexity of mechanism design," in *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence*, ser. UAI'02. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2002, pp. 103–110.

[15] D. Bergemann and S. Morris, "Robust mechanism design," *Econometrica*, vol. 73, no. 6, pp. 1771–1813, 2005.

[16] A. Penta, "Robust dynamic mechanism design," working paper, Apr 2013.

[17] P. Cramton, Y. Shoham, R. Steinberg, and V. Smith, *Combinatorial Auctions*. MIT Press, 2006.