Efficient Mechanisms with Dynamic Populations and Dynamic Types*

Ruggiero Cavallo
Yahoo! Research
New York, NY, USA
cavallo@post.harvard.edu

David C. Parkes
SEAS, Harvard University
Cambridge MA, USA
parkes@eecs.harvard.edu

Satinder Singh
Computer Science and Engineering
University of Michigan
Ann Arbor MI, USA
baveja@umich.edu

First version: June 5, 2007, Revised: December 19, 2008; July 7, 2009; September 30, 2009
This version: September 4, 2010

Abstract

We consider the truthful implementation of an efficient decision policy when agents have dynamic type and are periodically-inaccessible, with agents unable to report type information or make payments while inaccessible. This concept of inaccessibility includes arrival-departure dynamics as a special case. We generalize the dynamic-VCG (or pivot) mechanism [Bergemann and Välimäki, 2010] to this environment, emphasizing its position within a family of dynamic Groves mechanisms. In considering the special case of arrival-departure dynamics and dynamic type, the mechanism is within-period *ex post* incentive compatible so long as arrivals are independent of past arrivals, conditioned on actions of the center. For arrival-departure dynamics and static types, the mechanism is payoff equivalent upon arrival to the online-VCG mechanism [Parkes and Singh, 2003], serving to unify two previously disparate mechanisms and highlighting a tradeoff between *ex post* participation and *ex post* no-deficit properties.

*This paper subsumes a previous version entitled "Efficient Online Mechanisms for Persistent, Periodically Inaccessible Self-Interested Agents", dating from June, 2007. Thanks to the referees and editor of an earlier version for valuable feedback, seminar participants at Dagstuhl, Stonybrook, GAMES, Cornell, Yale, EPFL, NYU Stern, Aarhus, and Stanford, and Sven Seuken for useful comments.

1 Introduction

Many interesting problems of mechanism design are dynamic; for instance, a repeated allocation problem in which agents are learning their value for a resource, selling a time-sensitive good such as a theater ticket to impatient buyers that arrive at different times, or allocating a dynamic stream of user impressions to a shifting population of online advertisers. In designing incentive compatible mechanisms for dynamic and uncertain environments, existing models generally consider either a dynamic population with static, private information [Lavi and Nisan, 2004; Parkes and Singh, 2003] or a static population of agents with dynamic, privately revealed information [Athey and Segal, 2007; Cavallo *et al.*, 2006; Bergemann and Välimäki, 2010]. The former problems have been described as those of online auctions or online mechanism design (emphasizing the connection to *online* algorithm problems of computer science and operations research), while the latter have been described as problems of *dynamic* mechanism design (and allow for agents with uncertain local problems, including agent learning.)

In this paper, we present a unified model that covers both dynamic populations and dynamic types, by allowing agents to transition between states that may be inaccessible to a mechanism center. While inaccessible, an agent is unable to communicate with the center or receive payments. This model includes arrival-departure dynamics as a special case (where an agent becomes accessible for a contiguous number of periods), while extending efficient mechanism design to environments of practical interest in which agents may become disconnected from a mechanism because of faulty technology or reasons of limited attention or costly communication. For example, a bidder in an online auction becomes inaccessible when he is offline (but may remain interested in some goods in the market), a worker in an online labor market becomes inaccessible when he is focused on completing an assigned task (but may remain interested in being assigned a new task), and so forth. To motivate models of arrival-departure dynamics, consider a tourist who arrives in a new city and competes for theater tickets, a firm that decides on Monday to seek a temporary clerical assistant for a couple of days, or an R&D organization bidding for access to an on-demand computational 'cloud' facility.

Our results extend the dynamic-VCG (or pivot) mechanism [Bergemann and Välimäki, 2010] to this environment of dynamic populations and dynamic types. The generalized mechanism is within-period ex post incentive compatible, providing truthful reporting for accessible agents who form beliefs about the types of inaccessible agents that are consistent with the most recent reports of these agents. Specializing to arrival-departure dynamics and static types we also unify two disparate strands in the literature, by establishing payoff equivalence upon arrival between an instantiation of the dynamic-VCG mechanism [Bergemann and Välimäki, 2010] (originally described for a dynamic type, static population model) and the online-VCG mechanism [Parkes and Singh, 2003]. The two mechanisms are both within period *ex post* incentive compatible in this special case, but incomparable in terms of participation and no-deficit properties.

Our technical results emphasize that these mechanisms are all instances of a family of dynamic Groves mechanisms, insisting that the flow payoff to an agent in any period *in which it is accessible* is aligned with the total expected payoff of the system and some offset that is independent of its report in this or any future period. In addition to providing a simple explanation for incentive compatibility, familiar from standard Groves [Groves, 1973] intuitions, this yields a new proof of incentive compatibility of dynamic-VCG mechanisms. It is known that the incentive compatibility of the dynamic-VCG mechanism requires private values, and type transitions that are independent, conditioned on the actions of the center. Specialized to arrival-departure dynamics, we observe that this equates to arrivals that are independent of previous arrivals, conditioned on actions of the center. For example, the mechanisms described here fail to be incentive compatible if the probability that a high bid will be made on Tuesday depends on whether or not a high bid was made on Monday. We expect this observation to be important in practical applications.

1.1 Related work

We categorize the related work by the kind of dynamics considered. First, we review related work for a problem in which there is a persistent population of agents each with a dynamic type. This includes the general purpose, efficient mechanisms of Athey and Segal [2007], Bergemann and Välimäki [2010] and Cavallo et al. [2006], along with mechanisms that have been developed for special cases. Then we review related work for problems with a dynamic population of agents each with a static type. This includes the general purpose, efficient mechanism of Parkes and Singh [2003], as well as a number of mechanisms for special cases including both revenue-maximizing mechanisms and prior-free mechanisms that are analyzed within a worst-case rather than Bayesian framework.

Static population, dynamic type. Athey and Segal [2007] obtain a Bayes-Nash incentive compatible, efficient and budget-balanced mechanism for a persistent-population, dynamic type environment with private values and independent type transitions. The mechanism is *ex ante* individual rational, extending the expected externality mechanism [Arrow, 1979; d'Aspermont and Gérard-Varet, 1979]. The authors also provide sufficient conditions for *interim* participation in an infinite horizon setting with sufficiently patient agents. Bergemann and Välimäki [2010] introduce the dynamic-VCG (pivot) mechanism for the same environment, providing within-period *ex post* incentive compatibility and within-period *ex post* individual-rationality. The dynamic-VCG mechanism is *ex post* no deficit in economic environments without positive externalities (e.g., one-sided auctions and social-choice problems, but not double auctions or exchanges). The mechanism is unique when an *efficient exit* property is imposed, which requires that no transfers should occur once an agent's reports are no longer pivotal.[1] Applications are given to a schedul-

[1] The uniqueness result holds for static population, dynamic type environments but not for dynamic population, static type environments, in which incentive-compatibility constraints are only required upon arrival.

ing problem, and also to a problem with Bayesian learning by agents, modeled as a multi-armed bandits auction.[2] Cavallo et al. [2006] independently develop a variation on the dynamic-VCG mechanism for the same environment, modifying the team mechanism to provide within-period *ex post* incentive compatibility but achieving *ex ante* rather than (essentially) *interim* individual rationality. An application is given to a multi-armed bandits auction. Pavan et al. [2009] identify necessary and sufficient conditions for Bayes-Nash incentive compatibility, while requiring only *ex ante* participation. In addition to obtaining a general revenue-equivalence result, the analysis is applied to develop optimal dynamic mechanisms for problems in which agent type transitions are modeled as an auto-regressive process. Cavallo [2008] introduces a redistribution mechanism for multi-armed bandits auctions, returning payments to agents while retaining *ex post* no deficit and within-period *ex post* incentive compatibility, and in emphasizing the family of efficient, dynamic Groves mechanisms shows the revenue optimality of dynamic-VCG within this class.

An earlier literature developed dynamic mechanisms for persistent agents with time-separable types. For example, Atkeson and Lucas [1992] consider a continuum population in which agents receive new i.i.d. types each period, and characterize incentive-compatible distribution policies for a time-sensitive good. Athey et al. [2004] adopt a dynamic mechanism design approach in analyzing the equilibrium behavior of two competing firms, each with a private cost sampled i.i.d. in each period. Also related is the literature on dynamic contracting models, where the focus is on the role of commitment in limiting what a principal can achieve in problems with moral hazard; see Athey and Segal [2007] for a recent survey.

Dynamic population, static type. Parkes and Singh [2003] obtain the online-VCG mechanism for an environment with a dynamic population and static type, with a known probabilistic model of the arrival process and agents with private values on sequences of decisions.[3] The online-VCG mechanism allows for temporal strategies, with an agent "lurking" before reporting its type. Static type means that the value to an agent for all possible sequences of decisions by the center is known to the agent upon its arrival. The online-VCG mechanism is within-period *ex post* incentive compatible and efficient, collecting a single payment from an agent at, or subsequent to, its *commitment period*, which is the first period in which all decisions with respect to the agent's value are determined. The mechanism satisfies *ex post* participation and is *ex ante* no deficit in economic environments without positive

[2] Cremer et al. [2009] independently develop a special case of the Bergemann and Välimäki [2010] dynamic-VCG mechanism for an application with one-time type transitions, modeling costly information acquisition by agents. In an application in which each agent has instead a general, sequential process for costly value refinement, Cavallo and Parkes [2008] apply the dynamic-VCG mechanism, obtaining a reduction to a multi-armed bandits auction problem and extending to handle settings in which one agent can deliberate about the value of another agent.

[3] The authors establish that the online-VCG mechanism is Bayes-Nash incentive compatible. In the current paper, we emphasize that the online-VCG mechanism achieves this stronger, *ex post* incentive-compatibility property, and also formalize the requirement that arrivals be independent of past arrivals, conditioned on actions.

externalities.[4] Parkes et al. [2004] relax to ϵ-incentive compatibility and present an application to a scheduling problem. Mierendorff [2008] presents an application to bidders that have a value schedule for receiving an item in different time periods, obtaining a mechanism that is efficient and both *ex post* no deficit and *ex post* individual-rational by refactoring the payment flows across periods.[5]

Considering also revenue optimality, together with impatient buyers (that depart upon arrival and need an immediate decision), there are results selling identical items by a deadline to unit-demand buyers [Vulcano *et al.*, 2002] (see also [Gallien, 2003]), selling commonly ranked, distinct items to sell by a deadline to buyers with unit demand [Gershkov and Moldovanu, 2008a; 2008b], selling one unit of an identical good each period [Said, 2009], and for an adaptive mechanism that learns about the arrival process, under Bayesian and non-Bayesian learning paradigms [Gershkov and Moldovanu, 2008c; 2009].[6] For patient buyers, identical goods, and unit-demand, Pai and Vohra [2008] obtain regularity conditions for provably revenue-optimal auctions and Mierendorff [2009] studies the irregular case. For dynamic knapsack auctions, with a finite number of identical items to allocate by a deadline to buyers that demand a particular quantity, Dizdar et al. [2009] design efficient and revenue-optimal auctions for impatient buyers and Constantin and Parkes [2009] develop a tractable, heuristic approach for *patient* buyers.

A sequence of papers adopt *worst-case* rather than expected-case analysis for the design of online mechanisms. The objective is to develop mechanisms that can do well relative to what would be possible in a static problem with all type information available in the first period, and whatever the actual (dynamic) realization of agent types. Initiating this line of research, and adopting dominant-strategy incentive compatibility as a solution concept, Lavi and Nisan [2004] provide a worst-case analysis for an online auction in which a number of identical goods are sold by some deadline to agents with marginal-decreasing values for each additional good. Subsequent papers consider non-preemptive scheduling [Ng *et al.*, 2003], preemptive scheduling [Porter, 2004; Hajiaghayi *et al.*, 2005;

[4]The mechanism is developed in the context of a finite time horizon, but it is a simple matter to generalize it to an environment with an infinite time horizon and discounting, as we do in this paper.

[5]In a remarkable early contribution, Dolan [1978] developed an efficient mechanism for a scheduling problem with Poisson arrivals and agents with different delay costs. No consideration is given to temporal strategies, and agents are only able to misreport their cost of delay, not their time of arrival. The author characterizes the efficient policy and proposes to charge an agent the expected externality it imposes on the system upon its arrival. Dolan also observes that the dynamic mechanism will not be dominant strategy incentive compatible because it requires agreement about the probabilistic model of the arrival process, but offers no alternative equilibrium model.

[6]Also related is a literature that studies the problem of a monopolist selling multiple items of a durable good via posted price mechanisms to dynamic arrivals of impatient, unit-demand agents. For continuous time models, revenue-optimal price schedules are developed in a series of papers by Kincaid and Darling [1963], Gallego and Van Ryzin [1994] for an exponential demand model, and McAfee and te Velde [2008] for a Pareto demand model. Board [2008] develops revenue-optimal pricing schedules with buyers that arrive over time and discount the future, in a discrete time setting with time-varying demand.

Cole *et al.*, 2008], learning by the center via connections to the secretary problem [Hajiaghayi *et al.*, 2004], for an uncertain supply of expiring items to a static population [Mahdian and Saberi, 2006],[7] for general combinatorial valuations [Juda and Parkes, 2009], and in application to dynamic double auctions [Bredin *et al.*, 2007]. The scheduling problem has also been studied under alternative solution concepts, including set-Nash [Lavi and Nisan, 2005] and undominated strategies [Lavi and Segev, 2008].

2 The Model

We consider an environment with agents $I = \{1, \ldots, n\}$, and an action $a \in A$ selected by a center in each of a sequence of discrete time periods, perhaps infinite, from an action set A. Each agent $i \in I$ has a dynamic type $\theta_i = (s_i, \tau_i, r_i) \in \Theta$, where Θ is the set of possible types, that consists of *state*, $s_i \in S_i$ in state space S_i, a *stochastic transition function*, $\tau_i : S_i \times A \to S_i$, such that for all $s_i \in S_i$ and $a \in A$, $\sum_{s'_i \in S_i} Pr(\tau_i(s_i, a) = s'_i) = 1$ (for probability function Pr), and a *reward function*, $r_i : S_i \times A \to \Re$, which defines the value $r_i(s_i, a)$ to the agent for action a in state s_i. Whether or not an agent is accessible depends on its state. Let $\mathcal{A}^t(s_i) \in \{0, 1\}$ denote whether or not agent i is accessible given state s_i.

This is a private-values model, with reward $r_i(s_i, a)$ and next state $\tau_i(s_i, a)$, independent of the state of other agents, conditioned on action a.[8] An agent's transition and reward functions are invariant across time. It is the stochastic nature of transition function τ_i that makes this a dynamic type model, and requires receiving multiple reports from an agent (e.g., a report of the agent's current state in each period that it is accessible). For the special case in which τ_i is deterministic, each agent has a static type and its value is defined for all possible sequences of actions by reporting its type in the first period.

State $s_0 \in S_0$ and stochastic transition function $\tau_0 : S_0 \times A \to S_0$, both known to the center, allow the feasible actions in each period to depend on the history of actions (which can be captured through this state variable.) For example, in an auction this state includes information about which items have been allocated. A reward to the center (or designer) of $r_0(s_0, a)$ can be easily included and presents no additional technical difficulty. To keep our presentation simple, we ignore r_0 and adopt fixed action set A in what follows, but both can be easily included.

[7] Working in this model, Babaioff et al. [2009] provide a hybrid analysis that is worst-case with respect to agent valuations, but average-case with respect to a probabilistic model of supply and requires dominant-strategy incentive compatibility with regard to any possible supply realization, precluding the use of an online-VCG mechanism.

[8] Without inaccessibility, and thus precluding a setting with agent arrivals and departures, there are no additional technical difficulties in also allowing for serially correlated types, where transitions $\tau_i(s_i^k, a, \omega^k) \in S_i$ and rewards $r_i(s_i^k, a, \omega^k)$ also depend on an exogenous random process $\omega^0, \omega^1, \ldots$ observable to all agents. Serially correlated types are considered, for example, in, Athey and Segal [2007] and Cavallo [2008]. But this presents a problem with periodic inaccessibility, including arrival-departure dynamics, because it presents an information externality [Gershkov and Moldovanu, 2008c].

Let $\theta_0 \in \Theta_0$ define (s_0, τ_0), and $\boldsymbol{\Theta} = \Theta_0 \times \Theta^n$ denote the joint type space, with $\Theta^n = \prod_{i=1}^n \Theta$ and $\theta \in \boldsymbol{\Theta}$ a joint type profile. Similarly, let $S = S_0 \times S_1 \times \ldots \times S_n$ denote the joint state space. Let $\tau(s, a) = (\tau_0(s_0, a), \tau_1(s_1, a), \ldots, \tau_n(s_n, a)) \in S$ denote the joint transition function, and adopt $\theta_{-i} = (\theta_0, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots, \theta_n)$, analogously for s_{-i}. We assume discount factor $0 \leq \gamma \leq 1$, common to all agents and the designer, and let K denote a time horizon, perhaps infinite, with decision periods indexed as $t \in \{0, 1, \ldots, K\}$. For K infinite we require $\gamma < 1$.

2.1 Always accessible agents

If every agent is always accessible, the associated social planner's problem is a multi-agent Markov decision process (MDP). Each action taken by the center creates stochastic type transitions local to each agent, and generates a reward to each agent. A decision policy $\pi : S \to A$ specifies an action for every joint state profile $s \in S$.[9] Fixing transition functions τ, the expected discounted reward, or flow value, to agent i of policy π in state s^t is

$$V_i(s^t, \pi) = \mathbb{E}_{s^{t..K}} \left[\sum_{k=t}^{K} \gamma^{k-t} \; r_i^k(s_i^k, \pi(s^k)) \,\middle|\, s^t, \tau \right], \tag{1}$$

with $s^k = \tau(s^{k-1}, \pi(s^{k-1}))$ for $k > t$. Let $V(s^t, \pi) = \sum_{i \in I} V_i(s^t, \pi)$. An *efficient decision policy*, π^*, solves

$$\pi^* \in \arg\max_{\pi \in \Pi} V(s, \pi), \quad \forall s \in S, \tag{2}$$

where Π is the space of feasible policies. Let $V_{-i}(s^t, \pi) = \sum_{j \neq i} V_j(s^t, \pi)$ and let π^*_{-i} denote a policy that is efficient for agents other than i, i.e., $\pi^*_{-i} \in \arg\max_{\pi \in \Pi} V_{-i}(s, \pi), \forall s \in S$. The efficient policy depends on reward and transition functions, and we adopt notation π^*_θ (and similarly $\pi^*_{\theta_{-i}}$) when this context is not clear.

2.2 Periodically inaccessible agents

For environments with agents that are periodically inaccessible, there is a loss in efficiency to the social planner, who must make decisions based on probabilistic information about an inaccessible agent's state. For this, we define a *belief state* in period t, $\mathsf{s}_i^t \in \Delta(S_i)$, where $\Delta(S_i)$ is the set of probability distributions on agent i's states. This is the belief of an observer about the state of agent i, given knowledge of the state when the agent was last accessible and through Bayesian inference based on the subsequent sequence of actions and the agent's transition function. For an accessible agent, the belief state s_i^t assigns probability 1 to state s_i^t. Let $\mathsf{s}^t = (s_0^t, \mathsf{s}_1^t, \ldots, \mathsf{s}_n^t) \in \Delta(S) = S_0 \times \Delta(S_1) \times \ldots \times \Delta(S_n)$ denote the belief-state profile, where we include state $s_0^t \in S_0$, which is always known with certainty.

[9] In a finite time-horizon setting the decision policy should also depend on the current period t, but this is a detail we omit for notational simplicity.

The reward and transition functions can be extended in the natural way, with $\mathsf{r}_i(\mathsf{s}_i^t, a) = \mathbb{E}_{s_i^t}\left[r_i(s_i^t, a) \,\middle|\, \mathsf{s}_i^t\right]$, with $\mathsf{r}(\mathsf{s}^t, a) = \sum_i \mathsf{r}_i(\mathsf{s}_i^t, a)$, and transition function $\uptau_i : \Delta(S_i) \times A \to \Delta(S_i)$ defined to extend the stochastic transition function from states to belief states, so that $\uptau_i(\mathsf{s}_i^t, a)$ transitions to the belief state induced by Bayes rule and $\tau_i(s_i^t, a)$. Altogether, this defines a *belief type*, $\vartheta_i = (\mathsf{s}_i, \uptau_i, \mathsf{r}_i) \in \Theta$, with joint belief type space, $\boldsymbol{\Theta} = \Theta_0 \times \Theta^n$.[10] A decision policy is similarly extended to belief states, with $\pi : \Delta(S) \to A$, and we have flow value

$$V(\mathsf{s}^t, \pi) = \mathbb{E}_{\mathsf{s}^{t..K}}\left[\sum_{k=t}^{K} \gamma^{k-t}\ \mathsf{r}(\mathsf{s}^k, \pi(\mathsf{s}^k)) \,\middle|\, \mathsf{s}^t, \tau\right], \tag{3}$$

with analogous variants for V_i and V_{-i}. The efficient policy π^*, for periodically-inaccessible agents, solves $\pi^* \in \arg\max_{\pi \in \Pi} V(\mathsf{s}^t, \pi)$, $\forall \mathsf{s}^t \in \Delta(S)$, where Π denotes the space of feasible policies.[11] Let π^*_{-i} denote the policy that is efficient for the system of periodically-inaccessible agents $\neq i$. As in the always-accessible case, the efficient policy depends on the transition and reward functions, we typically write π^*_θ and $\pi^*_{\theta_{-i}}$ to make this dependence clear.

An agent's accessibility depends on its state, and thus on the actions taken by the center. Therefore, an efficient policy will inherently consider the value-of-information from an action that will make an agent accessible (e.g., allocating communication resources to an agent), thus revealing to the planner the state of a currently inaccessible agent.

3 The Generalized Dynamic-VCG Mechanism

A *dynamic mechanism*, $M = (\pi_\theta, x_\theta)$, is defined by a decision policy $\pi_\theta : \Delta(S) \to A$ and a transfer policy $x_\theta = \{x_{\theta,1}, \ldots, x_{\theta,n}\}$, with $x_{\theta,i} : \Delta(S) \to \Re$, for all $i \in I$. We focus on direct-revelation mechanisms.[12] Each agent can make a report about its type in each period in which it is accessible. In the usual case, it is only a change in state that an agent will wish to report but the formal set-up also allows an agent to report a change in transition or reward function. Note: in the special case of agents that are always accessible, these policies are defined on S instead of $\Delta(S)$.

In defining a mechanism's policies, π_θ, and x_θ, we make the type profile θ explicit because the type profile enters into the actions and payments of the mechanism in two different ways. To be concrete, consider a mechanism with an efficient decision policy. First, the transition and reward functions determine *which policy is efficient* because they define the MDP. Second, the state profile determines *the action selected for a given, efficient policy.*

[10] Note that $\mathsf{r}_i(\mathsf{s}_i^t, a)$ and $\uptau_i(\mathsf{s}_i^t, a)$ remain conditionally-independent of the belief states of other agents given the action of the center.

[11] For a computational approach to solve these belief-state MDP problems, see the survey of algorithms for Partially Observable MDPs in [Kaelbling *et al.*, 1996].

[12] Myerson [1986] and Green and Laffont [Green and Laffont, 1986] give a general revelation principle for dynamic communication games; see also [Athey and Segal, 2007; Parkes, 2007].

Type reports by agents and actions selected by the mechanism in each period are assumed to be public, although in a truthful equilibrium the strategies are invariant to this history. Thus, our results hold whether or not agents have information about this history. Let $h^t \in \mathcal{H}^t$ denote a sequence of type reports, actions and transfers, up to and including period t, from a set of possible histories $\mathcal{H}^t$. A strategy,

$$\sigma_i^t : \mathcal{H}^{t-1} \times \Theta \to \{\Theta \cup \phi\}, \tag{4}$$

defines the type report made by agent i given this history and given the agent's current type, where we use ϕ to indicate that the agent claims to be inaccessible.[13] For an inaccessible agent, so that $\mathcal{A}(s_i) = 0$, we require $\sigma_i(h^{t-1}, \theta_i) = \phi$.

Let $\sigma^t(h^{t-1}, \theta) = (\theta_0, \sigma_1^t(h^{t-1}, \theta_1), \ldots, \sigma_n^t(h^{t-1}, \theta_n)) \in \boldsymbol{\Theta}$, denote the reported type profile in period t. Let σ_i^* denote a truthful strategy, with

$$\sigma_i^*(h^{t-1}, \theta_i^t) = \begin{cases} \theta_i^t, & \text{, if } \mathcal{A}(s_i) = 1 \\ \phi & \text{otherwise,} \end{cases} \tag{5}$$

We first review the dynamic-VCG mechanism, which is defined for always accessible agents. Let $V_i(h^{t-1}, s^t, \pi_\theta, \sigma_i)$ denote the expected discounted value to agent i from the decisions under policy π_θ, after history h^{t-1}, in state s^t, given strategy σ_i, when the other agents are truthful. Let $X_i(h^{t-1}, s^t, \pi_\theta, x_\theta, \sigma_i)$ denote the analogous, expected discounted transfer to agent i. Agents have quasi-linear utility functions, so that an agent's expected discounted utility (or flow payoff) from state s^t under strategy σ_i is $V_i(h^{t-1}, s^t, \pi_\theta, \sigma_i) + X_i(h^{t-1}, s^t, \pi_\theta, x_\theta, \sigma_i)$.

Definition 1 (w.p. ex post IC)**.** *A dynamic mechanism $M = (\pi_\theta, x_\theta)$ in an environment with a fixed, accessible population and dynamic type, is within-period ex post incentive-compatible if, for every period t, for every agent $i \in I$, for every type profile $\theta^t \in \boldsymbol{\Theta}$, for every history h^{t-1}, and every $\sigma_i' \neq \sigma_i^*$, for truthful σ_i^*,*

$$\begin{aligned} V_i(h^{t-1}, s^t, \pi_\theta, \sigma_i^*) + & X_i(h^{t-1}, s^t, \pi_t h, x_\theta, \sigma_i^*) \geq \\ & V_i(h^{t-1}, s^t, \pi_\theta, x_\theta, \sigma_i') + X_i(h^{t-1}, s^t, \pi_\theta, x_\theta, \sigma_i'). \end{aligned} \tag{6}$$

In a within-period *ex post* incentive-compatible (wp-EPIC) mechanism [Bergemann and Välimäki, 2010; Athey and Segal, 2007], truthful revelation of state, transition and reward function is the best-response of an agent regardless of the current type profile, if all other agents are truthful in the current period and all future periods. This solution concept provides a strengthening of perfect Bayesian equilibrium [Fudenberg and Tirole, 1991], holding whatever the current beliefs of agents as long as agents play in equilibrium forward from the current period.

Definition 2 (Dynamic-VCG mechanism)**.** *A dynamic mechanism $M = (\pi_\theta^*, x_\theta)$ defined for an environment with a static population and dynamic type. In each*

[13] The strategy of agent i can also depend, in principle, on its own past types. We leave this out of the notation in the interest of simplicity. It does not change any of the results.

period t, given a type report $\theta_i^t = (s_i^t, \tau_i, r_i)$ (perhaps untruthful) from each agent i: action $a^t = \pi_\theta^*(s^t)$ is taken, where π_θ^* is the efficient policy given θ^t; and the following transfer is made to each agent i:

$$x_{\theta,i}^t(s^t) = r_{-i}(s^t, a^t) + \gamma \cdot \mathbb{E}_{s'}\Big[V_{-i}(s', \pi_{\theta_{-i}}^*) \,\Big|\, s' = \tau(s^t, a^t)\Big] - V_{-i}(s^t, \pi_{\theta_{-i}}^*). \qquad (7)$$

The second term is the expected optimal flow value to agents $\neq i$ forward from the next period given that action $a^t = \pi_\theta^*(s^t)$ is taken in the current period. Taken together, the transfer to agent i in each period is the flow marginal externality imposed on the other agents by its presence in the current period only.

Theorem 1. *[Bergemann and Välimäki, 2010] The dynamic-VCG mechanism is efficient and wp-EPIC in an environment with a fixed, accessible population, dynamic type, private values and independent type transitions conditioned on actions.*

We provide a new proof of this theorem in the appendix, emphasizing the positioning of the dynamic-VCG mechanism within a class of dynamic Groves mechanisms.

The dynamic-VCG mechanism is also within period *ex post* individual rational, with flow payoff $V_i(h^{t-1}, s^t, \pi_\theta^*, \sigma_i^*) - X_i(h^{t-1}, s^t, \pi_\theta^*, x_\theta, \sigma_i^*) \geq 0$ for all t, any type profile θ^t (and thus state s^t), and any history h^{t-1}. In particular, agent i's flow payoff in equilibrium forward from any type profile θ^t is equal to $V(s^t, \pi_\theta^*) - V_{-i}(s^t, \pi_{\theta_{-i}}^*)$, and is non-negative because of private values and since the feasible actions in a state are independent (conditioned on earlier actions) of the private types of agents.

For the general case, we define *consistent belief-type profile* $\breve{\vartheta}^t$ as the true type for accessible agents, with the belief state for an inaccessible agent as implied through Bayesian updates given the most recent state report (perhaps untruthful) received from the agent, and the transition and reward function for inaccessible agents equal to those in the most recent report (perhaps untruthful).

Given this, $V_i(h^{t-1}, \breve{\mathsf{s}}^t, \pi_\theta, \sigma_i)$ denotes the *expected discounted value* to an agent, i, accessible in period t, when it follows strategy σ_i, and adopts belief-state profile $\breve{\mathsf{s}}^t$ (from $\breve{\vartheta}^t$). We can similarly define the expected discounted transfer to agent i as $X_i(h^{t-1}, \breve{\mathsf{s}}^t, \pi_\theta, x_\theta, \sigma_i)$. For notational simplicity, and to be consistent with (3), we keep the dependence on the transition and reward profiles of agents hidden.

For belief states to be well-defined in every period, we assume in the sequel that all agents are accessible in the first period and insist that every agent makes an initial type report.[14]

Definition 3 (w.p. ex post IC with periodic inaccessibility)**.** *Consider an environment with periodic inaccessibility and dynamic types. Dynamic mechanism $M = (\pi_\theta, x_\theta)$, is within-period ex post incentive-compatible if, for every period t, for*

[14] This is without loss of generality, because we could instead adopt additional machinery, with a prior distribution assumed on initial agent types.

every accessible agent i, for every consistent belief-type profile $\breve{\vartheta}^t$ (and associated belief-state profile $\breve{\mathsf{s}}^t$), for every history h^{t-1}, and for every $\sigma'_i \neq \sigma^*_i$,

$$\begin{aligned} V_i(h^{t-1}, \breve{\mathsf{s}}^t, \pi_\theta, \sigma^*_i) + X_i(h^{t-1}, \breve{\mathsf{s}}^t, \pi_\theta, x_\theta, \sigma^*_i) \geq \\ V_i(h^{t-1}, \breve{\mathsf{s}}^t, \pi_\theta, \sigma'_i) + X_i(h^{t-1}, \breve{\mathsf{s}}^t, \pi_\theta, x_\theta, \sigma'_i) \end{aligned} \tag{8}$$

For an accessible agent, regardless of the current type of other accessible agents, and given consistent beliefs about the types of inaccessible agents, the agent's best-response is to be truthful given that other agents are also truthful forward from the current period.[15] We think of this as a communication-restricted within-period *ex post* Nash equilibrium, for a model in which agents 'know everything knowable' consistent with the communication constraints. No requirements are made of the mechanism in regard to incentive compatibility for an inaccessible agent because such an agent cannot make a type report in any case.

We are now ready to define the generalized dynamic-VCG mechanism.

Definition 4 (Generalized Dynamic-VCG Mechanism)**.** *A dynamic mechanism $M = (\pi^*_\theta, x^{\#}_\theta)$ defined for an environment with periodic inaccessibility and dynamic type. In each period t, given type reports from some subset of accessible agents: belief-type profile ϑ^t (and associated belief-state profile s^t) is updated based on these reports and according to Bayes rule for inaccessible agents; action $a^t = \pi^*_\theta(\mathsf{s}^t)$ is selected, where π^*_θ is the efficient policy given ϑ^t; and the following transfer is made to each agent i that makes a report:*

$$x^{\#}_{\theta,i}(\mathsf{s}^t) = \sum_{k=t-\delta(t)}^{t} \frac{x_{\theta,i}(\mathsf{s}^k)}{\gamma^{t-k}}, \text{ where} \tag{9}$$

$$x_{\theta,i}(\mathsf{s}^k) = \mathsf{r}_{-i}(\mathsf{s}^k, a^k) + \gamma \cdot \mathbb{E}_{\mathsf{s}'}\Big[V_{-i}(\mathsf{s}', \pi^*_{\theta_{-i}}) \,\Big|\, \mathsf{s}' = \tau(\mathsf{s}^k, a^k)\Big] - V_{-i}(\mathsf{s}^k, \pi^*_{\theta_{-i}}), \tag{10}$$

where $\pi^*_{\theta_{-i}}$ in (10) is efficient for agents $\neq i$ given belief-type profile ϑ^k_{-i}, and $\delta(t) \geq 0$ is the number of contiguous periods before period t that agent i claimed to be inaccessible.

Belief-state profile s' in the second term of (10) is obtained through Bayes rule given s^k and a^k, where the transition model and reward model adopted in $V_{-i}(\mathsf{s}', \pi^*_{\theta,-i})$ is given by ϑ^k_{-i} in period k. The two flow value terms to agents $\neq i$ in (10) are evaluated with respect to the transition and reward functions associated with belief-type profile ϑ^k_{-i} held by the center in period k.

[15] Truthful reporting is an equilibrium when accessible agents cannot know more about the state of inaccessible agents than implied by their most recent report. Suppose an agent j misreports its type in period t and goes inaccessible. Truthful reporting is not an equilibrium if one or more accessible agents has knowledge of the true type of j in period t, right before becoming inaccessible. On the other hand, as soon as agent j becomes accessible again (e.g., in some period t') then truthful reporting is an equilibrium even when all accessible agents know the current type of agent j, and even in period $t'+1$ following a misreport in period t' by agent j, because agent j will be truthful in period $t+1$ in equilibrium.

We say that an agent (perhaps inaccessible, or reported inaccessible) is *pivotal* in period t if its presence affects the action selected by the mechanism. In order to establish wp-EPIC we require the following assumption:

Assumption 1. *An agent that reports itself to be inaccessible in period t, and is pivotal in period t, must report itself as accessible in some future period.*

This ensures that agents cannot benefit from actions selected by the mechanism without eventually returning and making the *catch-up payment*, $x^{\#}_{\theta,i}(\mathsf{s}^t)$.[16]

Theorem 2. *The generalized dynamic-VCG mechanism is efficient and wp-EPIC for a fixed population of periodically-inaccessible agents given Assumption 1, private values, and independent type transitions conditioned on actions.*

To prove this result we establish that the mechanism is a generalized dynamic Groves mechanism, the class of which we now define formally. We defer the proof that the mechanism belongs to this class of mechanisms to the appendix. The essential idea is that the catch-up payment makes the mechanism payoff equivalent, for accessible agents, to a mechanism in which transfers $x_{\theta,i}(\mathsf{s}^t)$ are made in every period, whether or not an agent is accessible.

In defining the class of generalized dynamic Groves mechanisms, let $V_j(h^{t-1}, \breve{\mathsf{s}}^t, \pi_\theta, \sigma_i)$ denote an accessible agent's belief (given $\breve{\vartheta}^t$) about the expected discounted value to another agent j, perhaps inaccessible, when i adopts σ_i and the other agents are truthful. Define $V_{-i}(h^{t-1}, \breve{\mathsf{s}}^t, \pi_\theta, \sigma_i) = \sum_{j \neq i} V_j(h^{t-1}, \breve{\mathsf{s}}^t, \pi_\theta, \sigma_i)$ and $V(h^{t-1}, \breve{\mathsf{s}}^t, \pi_\theta, \sigma_i) = \sum_{j \in I} V_j(h^{t-1}, \breve{\mathsf{s}}^t, \pi_\theta, \sigma_i)$.

Definition 5 (generalized dynamic Groves mechanism)**.** *A generalized dynamic Groves mechanism $M = (\pi_\theta, x_\theta)$ in an environment with a fixed population of periodically-inaccessible agents and dynamic types,*

(i) receives type reports $\theta_i^t = (s_i^t, \tau_i, r_i)$ (perhaps untruthful) from some subset of accessible agents in each period t, and updates belief-type profile ϑ^t (and associated belief-state profile s^t) based on these reports and according to Bayes rule for inaccessible agents,

(ii) adopts a policy π_θ that is efficient with respect to the belief-type profile,

(iii) adopts a transfer policy x_θ so that, for every accessible agent i, for every consistent belief-type profile, $\breve{\vartheta}^t$ (and associated belief-state profile $\breve{\mathsf{s}}^t$), for every history h^{t-1} and every σ_i, and given that agents $\neq i$ are truthful in this and future periods, the expected discounted transfer is $V_{-i}(h^{t-1}, \breve{\mathsf{s}}^t, \pi_\theta, \sigma_i) - C_i(\breve{\mathsf{s}}^t_{-i})$, where $C_i(\breve{\mathsf{s}}^t_{-i})$ is independent of agent i's strategy in this period and forward.

[16] For strategic agents, one way to justify this assumption is to require an agent to post a bond that is returned over time, but only if the agent reports itself as accessible.

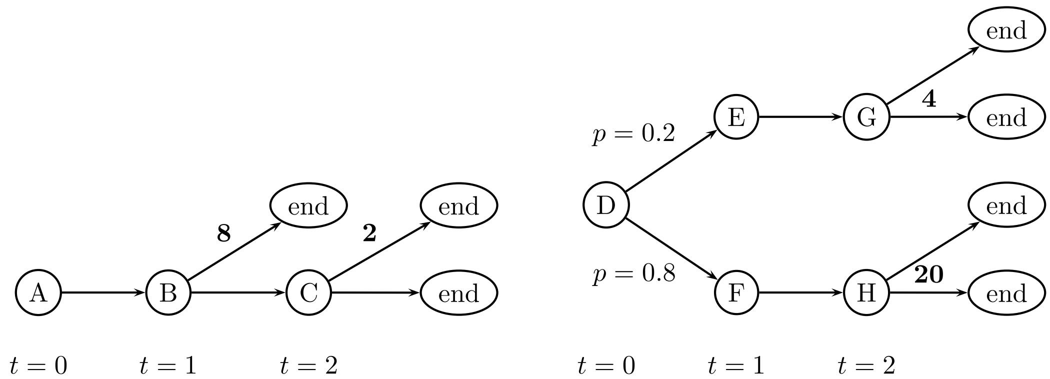

(a) Agent 1's type.

(b) Agent 2's type.

Figure 1: A dynamic allocation problem with two agents, a single item, and 3 time steps. Agent 1 has a static type (deterministic transitions) and agent 2 has a dynamic type. Rewards are shown in **bold** where non-zero, and correspond to the item being allocated to the agent.

Lemma 1. *A generalized dynamic Groves mechanism is efficient and wp-EPIC in an environment with a fixed population of periodically-inaccessible agents and dynamic types.*

Proof. Let π^*_θ denote an efficient policy on belief states. Fix agent i. Assume for contradiction that wp-EPIC fails. By the one-shot deviation principle,[17] there is some history h^{t-1}, some consistent belief type profile $\breve{\vartheta}^t$, and some strategy $\sigma_i \neq \sigma^*_i$ that provides more flow payoff to agent i than being truthful, while deviating only in period t (when agent i is accessible.) Assume that the other agents are truthful in period t and forward. By properties (i) and (ii), we have

$$V(h^{t-1}, \breve{\mathsf{s}}^t, \pi^*_\theta, \sigma_i) - C_i(\breve{\mathsf{s}}^t_{-i}) > V(\breve{\mathsf{s}}^t, \pi^*_\theta) - C_i(\breve{\mathsf{s}}^t_{-i}), \tag{11}$$

and so $V(h^{t-1}, \breve{\mathsf{s}}^t, \pi^*_\theta, \sigma_i) > V(\breve{\mathsf{s}}^t, \pi^*_\theta)$. Construct policy $\pi' : \Delta(S) \to A$ by setting $\pi'(\mathsf{s}) = \pi^*_\theta(\mathsf{s})$ in every belief state except for $\breve{\mathsf{s}}^t$, where π^*_θ is the efficient policy given $\breve{\vartheta}^t$ (which, recall, adopts agent i's true type whenever it is accessible). For belief state $\breve{\mathsf{s}}^t$, then $\pi'(\breve{\mathsf{s}}^t) = a'$, where action a' is selected by the mechanism given agent i's misreport, as prescribed by σ_i. By construction, $V(\breve{\mathsf{s}}^t, \pi') = V(h^{t-1}, \breve{\mathsf{s}}^t, \pi^*_\theta, \sigma_i)$ (where these flow values are all evaluated with respect to $\breve{\vartheta}^t$) and $V(\breve{\mathsf{s}}^t, \pi') > V(\breve{\mathsf{s}}^t, \pi^*_\theta)$ and a contradiction with the efficiency of π^*_θ. □

3.1 Example

To illustrate the generalized dynamic-VCG mechanism, we consider a simple example in which the center has a single item to allocate, there are two agents, and there is no discounting. As illustrated in Figure 1, agent 1's type is deterministic with value 8 for the item in period 1 and value 2 in period 2. Agent 2's type is dynamic, and is either value 4 or 20 for the item in period 2 depending on whether it is in state G or H. The system follows state dynamics $AD \rightarrow BE \rightarrow CG$ or $AD \rightarrow BF \rightarrow CH$. First assume that both agents are always accessible. In this case, the efficient policy allocates the item to agent 1 in state BE, to agent 2 in states $\{CG, CH\}$, and makes no allocation in states $\{AD, BF\}$. Consider state BE. If agent 2 is truthful, then agent 1 is allocated the item and receives transfer $0 + 0 - 4 = -4$ (i.e., makes payment 4). Now, if agent 2 lies and reports state F then no allocation is made and agent 2 receives transfer $0 + 2 - 8 = -6$. Agent 2 will actually transition to state G and the next joint state will be CG. Now, whatever the report of agent 2 (as long as it's > 2), it will be allocated the item and its transfer will be $0 + 0 - 2 = -2$. The actual value of the item to agent 2 is 4 and its total payoff is $4 - (6 + 2) = -4$, and worse than the zero payoff received when truthful. A similar analysis can be completed for state BF, and for agent 1, to confirm the incentive compatibility of the dynamic-VCG mechanism in this example.

Suppose now that agent 1 is accessible in every period with very high probability, while agent 2 is *inaccessible* in period 1 with high probability and becomes accessible in period 2 with high probability. Both agents are initially accessible in period 0. To gain intuition, first consider a naive mechanism with a policy that simply ignores inaccessible agents, picking actions as though such agents do not exist. The mechanism makes simple Groves payments: making a transfer to accessible agents in each period that is equal to the (reported) value received by other accessible agents. This is not incentive compatible. If agent 2 is inaccessible in period 1 then the mechanism will allocate to agent 1 (for a value of 8 to the agent.) But, if agent 1 also pretends to be inaccessible, the item will be allocated to agent 2 in period 2, for expected value $(0.2)4 + (0.8)20 = 16.8$. Agent 1's expected payoff for this strategy is 16.8 because it receives the value to agent 2 as a transfer. The problem arises because the policy is inefficient, ignoring the existence of agent 2, which is inaccessible in period 1.

Now consider the generalized dynamic-VCG mechanism. Let us suppose that agent 2 is accessible in period 1, and in state E. Suppose agent 1 is truthful. If agent 2 is truthful, then agent 1 would be allocated the item in period 1 and agent 2's payoff would be zero. But if agent 2 lies and pretends to be inaccessible in period 1, the efficient policy will delay making an allocation until period 2 because $8 < (0.2)4 + (0.8)20 = 16.8$ (ignoring the low probability of continued inaccessibility).

[17] For the one-shot deviation principle to hold it is sufficient that agents' beliefs off equilibrium satisfy *updating consistency*, which is a weakening of the requirements imposed by perfect Bayesian equilibrium [Fundenberg and Tirole, 1991; Hendon *et al.*, 1996; Perea, 2002].

Both agents' transfers in period 1 will be zero (agent 2's because it is inaccessible). Agent 2 can then report state G in period 2 and receive the item. But what about its transfer? It will receive a transfer of $-6-2=-8$ (i.e., make a payment of 8) when becoming accessible in period 2, which is more than its value for the item when its true state sequence is E and then G. (Note that if the agent could avoid its catch-up payment, it would have transfer -2 and obtain a net payoff of $4-2=2$, which would have been a worthwhile manipulation).

4 Special Case: Arrival and Departure Dynamics

We now consider a special case of periodic inaccessibility, in which each agent is first inaccessible, then accessible, and then inaccessible again, and where agents only have value for actions when they are accessible. This models an environment with arrival-departure dynamics and dynamic type. An arrival model provides the analogue to a belief state about the state of as yet unarrived agents, and the requirement that type transitions be independent, conditioned on actions, is interpreted here as a requirement that arrivals must be independent of past arrivals, conditioned on actions of the center.[18] A further specialization, to agents with a static type, provides a payoff equivalence (upon arrival) between the dynamic-VCG mechanism and the online-VCG mechanism [Parkes and Singh, 2003], unifying two disparate threads in the literature.

As a motivating example for arrival-departure dynamics, consider a tourist who arrives in New York City and is interested in buying theater tickets for multiple shows during a week; his type is dynamic while in town because he updates his value for different shows based on his experience attending earlier shows, or other unpredictable events such as the weather. However, once he leaves town at the end of the week, he is agnostic to subsequent actions of the mechanism.

To model arrival-departure dynamics within the periodically inaccessible model, we consider a potentially infinite set of agents $I=\{1,2,\dots,\infty\}$ (enough to model all possible arrivals), all initially inaccessible. It is convenient to introduce special states, `start` and `end`, to denote before arrival and after departure. Each agent i remains in its initial state `start` until the period in which it arrives, and $r_i(\mathtt{start},\mathtt{a})=\mathtt{0}$ for all $a \in A$, and transitions into some state $s_i \in S_i$ upon arrival, before finally transitioning into the `end` state upon departure, whereupon $r_i(\mathtt{end},\mathtt{a})=\mathtt{0}$ for all a.

A new state, $s_z^t \in S_z$, is introduced to state profile $s^t=(s_z^t, s_0^t, \{s_i^t : i \in I\})$, where $\Theta(s_z^t) \subset \Theta^{|I|}$ denotes the set of agent types that arrive in period t. In this way, state s_z models the arrival process, and has associated transition function $\tau_z : S_z \times A \to S_z$. Transition function τ_z is known to the mechanism while state s_z^t,

[18] For example, the arrival of a high bidder on Tuesday should not depend on whether or not a high bidder arrived on Monday. On the other hand, the arrival of a high bidder on Tuesday could depend on whether or not a resource was allocated on Monday, and this itself could depend on the types that arrive on Monday.

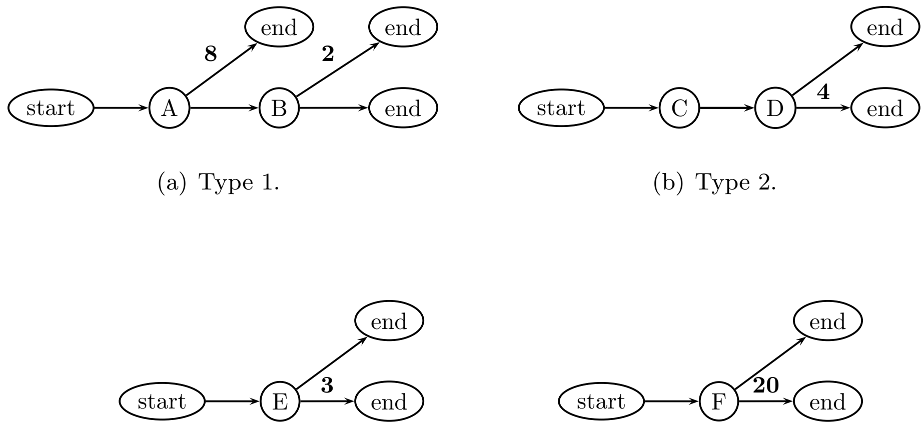

(a) Type 1.

(b) Type 2.

(c) Type 3.

(d) Type 4.

Figure 2: An example with arrival-departure dynamics, one item to allocate, two periods, and four possible types that an agent can have upon arrival. An agent with type 1 always arrives in period 1. One other agent arrives, with high probability, its arrival is in period 2 with type 4. With low probability, its arrival is in period 1 with either type 2 or type 3. Non-zero rewards are indicated in **bold**, and correspond to an allocation of the item to an agent with the associated type.

in any period t, and thus which agents, and with which types, arrive is unknown. Once an agent arrives, then its departure process conditioned on arrival is modeled by an eventual transition into the `end` state.

Note that agent i's strategy, σ_i, allows it to continue to claim to be inaccessible when it is accessible, and therefore includes the possibility of 'hiding' from the mechanism by delaying a report of its arrival, as well as claiming an early departure.

A natural mechanism to consider for this problem is to implement the efficient policy (based on reported types, and knowledge of the arrival process), and make the dynamic-VCG transfers (7) in periods where an agent reports itself to be accessible. But without a restriction on the arrival process, this (dynamic-VCG) mechanism is not incentive-compatible.

To see this, consider Figure 2, in which there are four possible agent arrival types, two periods, and there is one item to allocate. The arrival process is as follows. A first agent (agent 1) arrives in period 1 with (arrival) type 1, so that it is in state A upon arrival and it has value 8 and 2 for the item in periods 1 and 2 respectively. At most one additional agent arrives. With high probability, an agent with type 4 in period 2 (and in state F), with value 20 for the item. With low probability, either (i) an agent of type 2 arrives in period 1, in state C, and with value 4 for the item in period 2, or (ii) an agent of type 3 arrives in period 2, in state E, and with value 3 for the item. Suppose agent 1 is truthful, and that agent 2 has type 2 and therefore arrives in period 2. If truthful, the efficient policy will allocate to agent 1 in period 1 (knowing that an agent of type 4 will not arrive). But, if agent 2 hides and claims to arrive in period 2 and with type 3, the efficient action in period 1 is to hold the item, planning to allocate to a type 4 arrival in period 2. When agent 2 submits its report, pretending to be a type 3 agent in period 2, it receives the item and makes a transfer of 2 to the center, representing the externality imposed on agent 1, and achieve a payoff of 2.

The arrival process in the example fails to satisfy a necessary independence property. The probability that an agent transitions from `start` to F in period 2 (and thus a type 4 arrives) depends on whether or not an agent transitions from `start` to C in period 1 (and thus a type 2 arrives) because only one agent of type 2–4 can arrive. This independence property, where transitions are independent of states of other agents conditioned on actions, is critical for dynamic VCG to be incentive compatible.

To fix this, let $z^t = \Theta(s_z^t) \subset \Theta^{|I|}$ denote the set of types that arrive in period t, and insist on the following conditional independence property on the transition function associated with the arrival process:

Assumption 2. *An arrival process satisfies the conditional-independence on arrivals (CIA) property when*

$$\Pr(z^t \mid z^0, \ldots, z^{t-1}, a^0, \ldots, a^{t-1}) = \Pr(z^t \mid z'^0, \ldots, z'^{t-1}, a^0, \ldots, a^{t-1}), \tag{12}$$

for all $z^0, \ldots, z^{t-1}$*, all* $z'^0, \ldots, z'^{t-1}$*, all* $a^0, \ldots, a^{t-1}$*, and all* z^t*.*

This excludes the kind of arrival dynamics in the example. Still,whether or not a high type arrives in some period can depend on whether or not an item has been allocated, and this can in turn depend on the earlier types of agents.

The social planner's problem can now adopt state $(z^t, s_0^t, \{s_i^t : i \in I\}) \in S$, with $z^t \subset \Theta^{|I|}$ to represent the types of new arrivals in period t. The CIA property requires that transition $z^{t+1} = \tau_z(z^t, a)$ is independent of the current arrival state z^t, conditioned on action a.

The efficient policy is denoted π^*_θ, and depends on the arrival process τ_z as well as the transition and reward types of accessible agents. The efficient policy, $\pi^*_{\theta_{-i}}$, without agent i is identical to the efficient policy with agent i except when agent i is present (i.e., arrived and not yet departed). To see this, note that: (a) excluding an agent that has yet to arrive has no effect on the efficient policy because the arrival process determines the impact of possible arrivals on the policy; and (b) excluding an agent that has departed has no effect because a departed agent has no ongoing value for actions, and does not affect the probability of future arrivals by CIA. Moreover, all agent types satisfy Assumption 1.

Based on this observation, that an agent is only pivotal when present in the system, the catch-up payments in the generalized dynamic-VCG mechanism are zero. From this, we have a simplified mechanism, that is exactly the natural dynamic-VCG mechanism suggested earlier, in which the efficient policy given knowledge of the arrival process is adopted by the center, and the transfers (7) are made to each agent that makes a report in a period.

Theorem 3. *The dynamic-VCG mechanism for a population with arrival-departure dynamics and dynamic types, is efficient and wp-EPIC given private values, independent type transitions, the CIA property, and a center with a correct model of agent arrivals.*

The proof is deferred until the appendix. The crux of the argument is to recognize that the catch-up payments are always zero and therefore the transfers are equivalent to those that would be made to an agent if they could be made in every period. By wp-EPIC, truthful reporting is an equilibrium whatever the current type profile and whatever the arrival process, as long as other agents report their true types upon arrival, the center's model of the arrival process is correct, and it is common knowledge that the center's model is correct.[19] Note that the arrival model itself does not need to be known to agents.

4.1 Static Types and Arrival-Departure Dynamics

A special case with considerable practical importance is that of a dynamic population of agents with static types, that is an agent for whom the value for any sequence of actions is fixed, and known to the agent upon arrival. For example, an agent might have an invariant, time-separable valuation function or assign a value to a particular bundle of goods if received by a deadline. A static type has a deterministic transition function, $\tau_i : S_i \times A \rightarrow S_i$. With this, a single report of type $\theta_i^t = (s_i^t, \tau_i, r_i)$ upon arrival provides knowledge to an observer (such as the center) of the agent's state, and thus value for actions, in all future periods. Figure 3 provides a simple illustration of simple static types.

The dynamic-VCG mechanism is wp-EPIC and efficient in this domain, under the CIA property. Our interest in this section is to establish the same result for the online-VCG mechanism. The online-VCG mechanism is also wp-EPIC and efficient, but offers a different compromise between individual-rationality and no-deficit properties than the dynamic-VCG mechanism for dynamic, multi-unit auctions with values that are complements. These distinct economic properties may be important in some practical settings.

With static types, it is sufficient to consider report-once mechanisms, in which an agent can make only one report of its type. Given this, we can further specialize wp-EPIC so that it only needs to hold up until the period in which an agent reports its arrival (and not subsequent periods.) We need to define a *partially-truthful type profile* $\breve{\theta}^t$ (and associated state profile $\breve{s}^t$), which combines earlier reports (perhaps untruthful) received from agents with the true type of agents that either arrive in

[19]The 'communication-restricted' variation on wp-EPIC for the generalized dynamic-VCG mechanism in Definition 3 is irrelevant here because agents that are inaccessible are never pivotal.

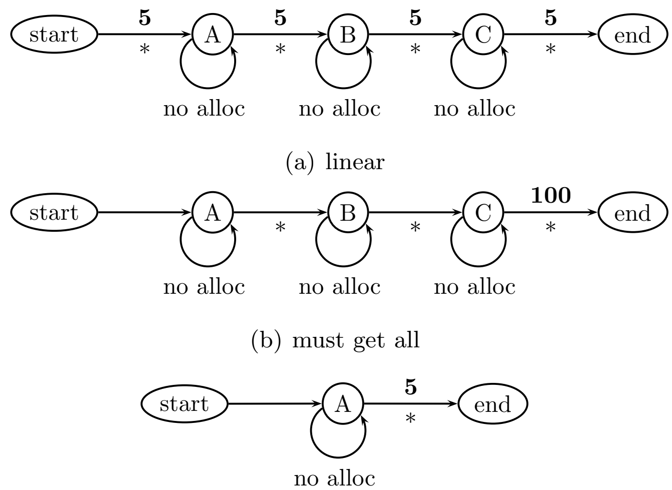

(a) linear

(b) must get all

(c) unit demand

Figure 3: Example valuations for agents with a static type in a dynamic, multi-unit allocation problem. An allocation of one unit is indicated by '*' with non-zero values in **bold**. The agent valuations are: (a) additive, with value 5 for each of up to three units of the item; (b) complements, with value 100 for three units; and (c) value 5 for a single unit.

period t or are present but did not yet submit a report. An agent is *active* in period t if the agent has arrived and not yet departed.

Definition 6 (w.p. ex post IC for arrival-departure dynamics and static type)**.** *A report-once dynamic mechanism, $M = (\pi_\theta, x_\theta)$, is within-period ex post incentive-compatible in a population with arrival-departure dynamics and static types if, for any period t, any agent i active in t but yet to report, any partially-truthful type $\breve{\theta}^t$ (with associated state $\breve{s}^t$), any history h^{t-1}, and for all $\sigma_i' \neq \sigma_i^*$,*

$$V_i(h^{t-1}, \breve{s}^t, \pi_\theta, \sigma_i^*) + X_i(h^{t-1}, \breve{s}^t, \pi_\theta, x_\theta, \sigma_i^*) \geq V_i(h^{t-1}, \breve{s}^t, \pi_\theta, \sigma_i') + X_i(h^{t-1}, \breve{s}^t, \pi_\theta, x_\theta, \sigma_i'), \tag{13}$$

where σ_i^* is truthful and immediate reporting upon arrival.

By (13), truthful reporting is at least as good as any other strategy σ_i', including delaying an agent's report by one or more periods. In the same way, a dynamic Groves mechanism for this problem only needs to align an agent's flow transfer up until the agent has submitted a type report:

Lemma 2. *A report-once dynamic mechanism (π_θ, x_θ) in an environment with arrival-departure dynamics and static types is efficient and wp-EPIC if:*

(i) policy π_θ is efficient for the correct arrival process, given agents' reported types, and

(ii) the expected discounted transfer to any agent i, active in period t and still to report (including a new arrival), given any partially-truthful type profile $\breve{\theta}^t$ (and associated state profile $\breve{s}^t$), history h^{t-1}, strategy σ_i, and given that agents $\neq i$ (still to report) are truthful going forward, is $V_{-i}(h^{t-1}, \breve{s}^t, \pi_\theta, \sigma_i) - C_i(\breve{s}^t_{-i})$, where $C_i(\breve{s}^t_{-i})$ is independent of agent i's strategy in this period and forward.

The proof of this lemma is omitted, but follows the same outline as earlier.[20] We now review the online-VCG mechanism.[21]

Definition 7 (online-VCG mechanism)**.** *In the online-VCG mechanism, $M = (\pi^*_\theta, x_\theta)$, for arrival-departure dynamics and static types, π^*_θ is the efficient policy given the known arrival process, and each agent can make a single report (perhaps untruthfully) about its type $\theta^t_i = (s^t_i, \tau_i, r_i)$ in some period t (its reported arrival). In every period t, given reported type θ^t, the mechanism selects action $a^t = \pi^*_\theta(s^t)$. In every period t between an agent i's reported arrival and departure, the mechanism makes transfer to i:*

$$x_{\theta,i}(s^t) = \begin{cases} -r_i(s^t, a^t) + \left[V(s^t, \pi^*_\theta) - V_{-i}(s^t, \pi^*_{\theta_{-i}})\right], & \textit{if arrived in period } t \\ -r_i(s^t, a^t) & \textit{otherwise,} \end{cases} \tag{14}$$

Theorem 4. *The online-VCG mechanism is efficient and wp-EPIC in an environment with arrival-departure dynamics, static types, private values, the CIA property, and a center with a correct model of the arrival process.*

The key to establishing the wp-EPIC of the online-VCG mechanism is that it is sufficient to show that the transfers satisfy property (ii) of Lemma 2 in the single period in which an agent could report its type. The online-VCG mechanism does not align incentives in the same way in subsequent periods, but at that point it matters not because the agent's type report is committed. The proof is deferred until the appendix.

Both the online-VCG and dynamic-VCG mechanisms are efficient and wp-EPIC with arrival-departure dynamics and static type. But they are different mechanisms, distributing transfers differently across states. This affects their properties in regard to individual rationality and no-deficit. The online-VCG mechanism, but not the dynamic-VCG mechanism, is *ex post* individual rational, with the payoff to an agent

[20] Suppose that the mechanism is not wp-EPIC, and establish a contradiction with the efficiency of the policy by showing that the center could instead adopt the actions taken by the policy given agent i's misreports and obtain a policy with greater total expected discounted value.

[21] The online-VCG mechanism was originally proposed by Parkes and Singh [2003] in a setting without discounting, with a transfer that is zero except upon departure, when an agent's transfer is $-\sum_{k=\underline{\ell}}^{\overline{\ell}} r_i(s^k_i, \pi^*_\theta(s^k)) + V(s^{\underline{\ell}}, \pi^*_\theta) - V_{\theta_{-i}}(s^{\overline{\ell}}, \pi^*_{\theta_{-i}})$, where $\underline{\ell}$ and $\overline{\ell}$ are the reported arrival and departure period of the agent. The online-VCG mechanism presented here is a simple variant that provide the analogous payoffs in a setting with discounting by refactoring transfers across periods.

	Agent 1		Agent 3	
Online VCG	Win	0.145	Win	$v_3 - 0.7$
	Lose	0.145	Lose	0
Dynamic VCG	Win	$0.6 - v_3$	Win	$v_3 - 0.7$
	Lose	-0.1	Lose	0

Table 1: Case analysis of payoff to agents 1 and 3 in the online-VCG and dynamic-VCG mechanisms in Example 1.

equal to the marginal expected product contributed by the agent at its arrival. The dynamic-VCG mechanism is within-period ex post individual rational, meaning it guarantees *expected* utility going forward is positive, but an agent may end up with negative utility depending on how random transitions are realized.

Example 1. *Consider a problem with two periods, one item to allocate, and three agents. In period 1, there is an agent with value 0.7 and arrival-departure (1,2) so it is patient. Another agent has value 0.6 and arrival-departure (1,1) so it is impatient. In period 2, there is an agent with value* $v_3 \sim U(0,1)$. *Suppose there is no discounting* ($\gamma = 1$). *The efficient policy is to wait until period 2 and then allocate to agent 1 if* $v_3 \leq 0.7$, *and to agent 3 otherwise.*

(a) Online-VCG. We have $V(\theta^1, \pi^*) = (0.7)(0.7) + (0.3)(0.85) = 0.745$ *and* $V_{-1}(\theta^1, \pi^*) = 0.6$. *If agent 1 wins, then its total transfer is* $-0.7 + (0.745 - 0.6) = -0.555$. *If agent 1 loses, then its total transfer is* $0 + (0.745 - 0.6) = 0.145$. *Agent 2 makes no transfer because it is not pivotal. Agent 3 receives no transfer if it loses, and its transfer is* $-v_3 + (v_3 - 0.7) = -0.7$ *if it wins.*

(b) Dynamic-VCG. We have $\mathbb{E}_{\theta'}[V_{-1}(\theta', \pi^*_{-1}) \,|\, \textit{don't allocate}] = 0.5$ *(the expected value to other agents in period 2 given that the item is not allocated in period 1), and* $V_{-1}(\theta^1, \pi^*_{-1}) = 0.6$. *In period 1, agent 1's transfer is* $0 + (0.5 - 0.6) = -0.1$. *In period 2, if agent 1 wins (i.e., if* $v_3 \leq 0.7$*) then its transfer is* $0 + (0 - v_3)$*; otherwise, if agent 1 loses then its transfer is 0. Agent 2 makes no transfer because it is not pivotal. Agent 3's transfer is 0 in period 1 because it has not arrived. If agent 3 wins, then it receives transfer* $0 + (0 - 0.7) = -0.7$ *and has no transfer if it loses.*

The payoff to agents 1 and 3 is tabulated in Table 1. Although the expected payoff is the same in both mechanisms, in dynamic-VCG agent 1's payoff is negative when it loses, or even if it wins and value $v_3 \in (0.6, 0.7]$. *The* ex post *payoff in the online-VCG mechanism is always non-negative.*

On the other hand, the dynamic-VCG mechanism but not the online-VCG mechanism is *ex post* no-deficit in economic environments without positive externalities (e.g., social choice and one-sided auction problems). The online-VCG mechanism satisfies *ex ante* no-deficit in these same environments.

Example 2. *Consider a problem with two periods and two units of an item to allocate. In period 1, there is a patient agent with value 100 for both units together*

	Period 2 arrivals	Agent 1 transfer	Other agents	Total transfer
Online VCG	0	$(-2\epsilon+\epsilon^2)100$	0	$(-2\epsilon+\epsilon^2)100$
	1	$100+(-2\epsilon+\epsilon^2)100$	-100	$(-2\epsilon+\epsilon^2)100$
	2	$100+(-2\epsilon+\epsilon^2)100$	0	$100+(-2\epsilon+\epsilon^2)100$
Dynamic VCG	0	0	0	0
	1	0	-100	-100
	2	0	0	0

Table 2: Case analysis of the transfers in the online-VCG and dynamic-VCG mechanisms in Example 2.

and arrival-departure (1,2). In period 2, up to two agents might arrive, each with low probability $\epsilon > 0$ of arriving and with value 150 for one unit. There is no discounting. The efficient policy is to wait until period 2 and allocate to agent 1 unless one or both of the high-value agents arrive, in which case these agent(s) are allocated instead.

(a) Online-VCG. We have $V(\theta^1,\pi^*) = (1-2\epsilon+\epsilon^2)(100)+2\epsilon(1-\epsilon)(150)+\epsilon^2(300)$ and $V_{-1}(\theta^1,\pi^*) = 2\epsilon(1-\epsilon)(150)+\epsilon^2(300)$. If agent 1 wins, then its total transfer is $-100+((1-2\epsilon+\epsilon^2)(100)+2\epsilon(1-\epsilon)(150)+\epsilon^2(300))-(2\epsilon(1-\epsilon)(150)+\epsilon^2(300)) = (-2\epsilon+\epsilon^2)(100)$. If agent 1 loses, then its total transfer is $100+(-2\epsilon+\epsilon^2)(100)$. If one agent arrives in period 2, it wins and its transfer is $-150+(150-100) = -100$. If two agents arrive in period 2, they each win and have transfer $-150+(300-150) = 0$.

(b) Dynamic-VCG. If agent 1 wins or loses, its transfer is 0. If one agent arrives in period 2, then it wins and its transfer is $0+(0-100) = -100$. If two agents arrive in period 2, then they each win and have transfer $150+(0-150) = 0$.

The transfers that occur in each mechanism, as they depend on the number of agents that arrive in period 2, are detailed in Table 2. Whereas the total expected transfer is the same in the two mechanisms, the online-VCG mechanism incurs a deficit of 100 (as $\epsilon \to 0$) when 2 agents arrive in period 2, while the dynamic-VCG mechanism always runs without a deficit.

5 Closing Remarks

In this paper, we have extended the dynamic-VCG mechanism [Bergemann and Välimäki, 2010] to environments in which agents are periodically-inaccessible, and used this extension to derive a mechanism that is wp-EPIC and efficient for problems with dynamic populations and dynamic agent type. For our results we require private values and type transitions that are conditionally independent of the types of other agents, when conditioned on actions by the center. From this, agent arrivals must be conditionally-independent of earlier arrivals, when conditioned on actions of the center. For the practically importance case of dynamic population and static type (capturing dynamic auctions, for example), we showed that the dynamic-VCG mechanism is payoff equivalent upon arrival to the online-VCG mechanism [Parkes and Singh, 2003], while differing in the exact timing of payments, with the former providing ex post no-deficit and the latter ex post individual rationality.

There are many interesting directions in developing a theory that can be instructive in the design of practical mechanisms for dynamic environments. Amongst the most exciting are (a) a theory for when dynamic VCG mechanisms satisfy core properties in dynamic combinatorial auctions to parallel results for the static VCG mechanism [Bikhchandani and Ostroy, 2002; Ausubel and Milgrom, 2006], (b) developing indirect mechanisms with efficient elicitation properties (see [Said, 2009] for initial work on this), (c) developing dynamic auctions that satisfy both ex post no-deficit and ex post individual-rationality in combinatorial domains, (d) developing representation languages to facilitate compact and expressive representations of dynamic agent types, (e) developing scalable computational methods to implement the mechanisms in real-world domains.

References

[Arrow, 1979] Kenneth J. Arrow. The property rights doctrine and demand revelation under incomplete information. In M. Boskin, editor, *Economics and Human Welfare*. Academic Press, 1979.

[Athey and Segal, 2007] Susan Athey and Ilya Segal. An efficient dynamic mechanism. Working paper, Stanford University. Earlier version circulated in 2004., 2007.

[Athey *et al.*, 2004] Susan Athey, Kyle Bagwell, and Chris Sanchirico. Collusion and price rigidity. *The Review of Economic Studies*, 71:317–349, 2004.

[Atkeson and Lucas, 1992] Andrew Atkeson and Robert E Lucas. On efficient distribution with private information. *The Review of Economic Studies*, 59:427–453, 1992.

[Ausubel and Milgrom, 2006] Lawrence M. Ausubel and Paul Milgrom. The lovely but lonely vickrey auction. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial Auctions*, chapter 1, pages 17–40. MIT Press, 2006.

[Babaioff *et al.*, 2009] Moshe Babaioff, Liad Blumrosen, and Aaron Roth. Auctions with online supply. In *Fifth Workshop on Ad Auctions*, 2009.

[Bergemann and Välimäki, 2010] Dirk Bergemann and Juuso Välimäki. The dynamic pivot mechanism. *Econometrica*, 78(2):771–789, 2010. Earlier version circulated as “Efficient Dynamic Auctions”, 2006.

[Bikhchandani and Ostroy, 2002] Sushil Bikhchandani and Joseph M Ostroy. The Package Assignment Model. *Journal of Economic Theory*, 107(2):377–406, 2002.

[Board, 2008] Simon Board. Durable-goods monopoly with varying demand. *The Review of Economic Studies*, 75:391–413, 2008.

[Bredin *et al.*, 2007] Jonathan Bredin, David C. Parkes, and Quang Duong. Chain: A dynamic double auction framework for matching patient agents. *Journal of Artificial Intelligence Research*, 30:133–179, 2007. Earlier version appeared in Proc. UAI'05.

[Cavallo and Parkes, 2008] Ruggiero Cavallo and David C. Parkes. Efficient metadeliberation auctions. In *Proc. 23rd AAAI Conference on Artificial Intelligence (AAAI'08)*, pages 50–56, Chicago, IL, 2008.

[Cavallo *et al.*, 2006] Ruggiero Cavallo, David C. Parkes, and Satinder Singh. Optimal coordinated planning amongst self-interested agents with private state. In *Proceedings of the Twenty-second Annual Conference on Uncertainty in Artificial Intelligence (UAI'06)*, pages 55–62, 2006.

[Cavallo, 2008] Ruggiero Cavallo. Efficiency and redistribution in dynamic mechanism design. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC-08)*, 2008.

[Cole *et al.*, 2008] R. Cole, S. Dobzinski, and L. Fleischer. Prompt mechanisms for online auctions. In *Symposium on Algorithmic Game Theory (SAGT)*, pages 170–181, 2008.

[Constantin and Parkes, 2009] Florin Constantin and David C. Parkes. Self-Correcting Sampling-Based Dynamic Multi-Unit Auctions. In *10th ACM Electronic Commerce Conference (EC'09)*, 2009.

[Crémer *et al.*, 2009] Jacques Crémer, Yossi Spiegel, and Charles Z. Zheng. Auctions with costly information acquisition. *Economic Theory*, 38:41–72, 2009.

[d'Aspermont and Gérard-Varet, 1979] C. d'Aspermont and L.A. Gérard-Varet. Incentives and incomplete information. *Journal of Public Economics*, 11:25–45, 1979.

[Dizdar *et al.*, 2009] Deniz Dizdar, Alex Gershkov, and Benny Moldovanu. Revenue maximization in the dynamic knapsack problem. Working paper, University of Bonn, 2009.

[Dolan, 1978] Robert J Dolan. Incentive mechainsms for priority queuing problem. *Bell Journal of Economics*, 9:421–436, 1978.

[Fundenberg and Tirole, 1991] Drew Fundenberg and Jean Tirole. Perfect bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory*, 53(2):236–260, 1991.

[Gallego and van Ryzin, 1994] Guillermo Gallego and Garrett van Ryzin. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, 40:999–1020, 1994.

[Gallien, 2003] Jeremie Gallien. Dynamic mechanism design for online commerce. *Operations Research*, 54:291–310, 2003.

[Gershkov and Moldovanu, 2008a] Alex Gershkov and Benny Moldovanu. Dynamic revenue maximization with heterogeneous objects: A mechanism design approach. *American Economic Journal: Microeconomics*, 2008. To appear.

[Gershkov and Moldovanu, 2008b] Alex Gershkov and Benny Moldovanu. Efficient sequential assignment with incomplete information. *Games and Economic Behavior*, 2008. To appear.

[Gershkov and Moldovanu, 2008c] Alex Gershkov and Benny Moldovanu. Learning about the future and dynamic efficiency. *American Economic Review*, 2008. To appear.

[Gershkov and Moldovanu, 2009] Alex Gershkov and Benny Moldovanu. Optimal search, learning and implementation. Working paper, University of Bonn, 2009.

[Green and Laffont, 1986] Jerry R. Green and Jean-Jacques Laffont. Partially verifiable information and mechanism design. *Review of Economic Studies*, 53(3):447–456, 1986.

[Groves, 1973] Theodore Groves. Incentives in Teams. *Econometrica*, 41:617–631, 1973.

[Hajiaghayi *et al.*, 2004] Mohammad T. Hajiaghayi, Robert Kleinberg, and David C. Parkes. Adaptive limited-supply online auctions. In *Proc. ACM Conf. on Electronic Commerce*, pages 71–80, 2004.

[Hajiaghayi *et al.*, 2005] Mohammad T. Hajiaghayi, Robert Kleinberg, Mohammad Mahdian, and David C. Parkes. Online auctions with re-usable goods. In *Proc. ACM Conf. on Electronic Commerce*, pages 165–174, 2005.

[Hendon *et al.*, 1996] E. Hendon, J. Jacobsen, and B. Sloth. The one-shot-deviation principle for sequential rationality. *Games Econ. Behav.*, 12:274–282, 1996.

[Juda and Parkes, 2009] Adam I. Juda and David C. Parkes. An options-based solution to the sequential auction problem. *Artificial Intelligence*, 173:876–899, 2009. Early version at AMEC workshop (2004) and ACMEC (2006).

[Kaelbling *et al.*, 1996] Leslie Pack Kaelbling, Michael L. Littman, and Andrew W. Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996.

[Kincaid and Darling, 1963] Wilfred M. Kincaid and Donald A. Darling. An inventory pricing problem. *Journal of Mathematical Analysis and Applications*, 7:183–208, 1963.

[Lavi and Nisan, 2004] Ron Lavi and Noam Nisan. Competitive analysis of incentive compatible on-line auctions. *Theoretical Computer Science*, 310:159–180, 2004. Earlier version in ACMEC 2000.

[Lavi and Nisan, 2005] Ron Lavi and Noam Nisan. Online ascending auctions for gradually expiring goods. In *Proc. of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1146–1155, 2005.

[Lavi and Segev, 2008] R. Lavi and E. Segev. Efficiency levels in sequential auctions with dynamic arrivals. Working paper, The Technion, 2008.

[Mahdian and Saberi, 2006] M. Mahdian and A. Saberi. Multi-unit auctions with unknown supply. In *Proc. ACM conference on Electronic Commerce (EC)*, pages 243–249, 2006.

[McAfee and te Velde, 2008] R. Preston McAfee and Vera te Velde. Dynamic pricing with constant demand elasticity. *Production and Operations Management*, 17:432–438, 2008.

[Mierendorff, 2008] Konrad Mierendorff. Efficient intertemporal auction. Working paper, University of Bonn, 2008.

[Mierendorff, 2009] Konrad Mierendorff. Optimal dynamic mechanism design with deadlines. Working paper, University of Bonn, 2009.

[Myerson, 1986] Roger Myerson. Multistage games with communication. *Econometrica*, 54(2):323–358, 1986.

[Ng *et al.*, 2003] Chaki Ng, David C. Parkes, and Margo Seltzer. Virtual Worlds: Fast and Strategyproof Auctions for Dynamic Resource Allocation. In *Proc. Fourth ACM Conf. on Electronic Commerce (EC'03)*, pages 238–239, 2003.

[Pai and Vohra, 2008] Mallesh Pai and Rakesh Vohra. Optimal dynamic auctions. Working paper, MEDS Department, Kellogg School of Management, 2008.

[Parkes and Singh, 2003] David C. Parkes and Satinder Singh. An MDP-based approach to Online Mechanism Design. In *Proc. 17th Annual Conf. on Neural Information Processing Systems (NIPS'03)*, 2003.

[Parkes *et al.*, 2004] David C. Parkes, Satinder Singh, and Dimah Yanovsky. Approximately efficient online mechanism design. In *Proc. 18th Annual Conf. on Neural Information Processing Systems (NIPS'04)*, 2004.

[Parkes, 2007] David C Parkes. On-line mechanisms. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani, editors, *Algorithmic Game Theory*, chapter 16. Cambridge University Press, 2007.

[Pavan *et al.*, 2009] Alessandro Pavan, Ilya Segal, and Jusso Toikka. Dynamic mechanism design: Incentive compatibility, profit maximization and information disclosure. Working paper, Northwestern University, 2009.

[Perea, 2002] Andrés Perea. A note on the one-deviation property in extensive form games. *Games and Economic Behavior*, 40(2):322–338, 2002.

[Porter, 2004] Ryan Porter. Mechanism design for online real-time scheduling. In *Proc. ACM Conf. on Electronic Commerce (EC'04)*, pages 61–70, 2004.

[Said, 2009] Maher Said. Auctions with dynamic populations: Efficiency and revenue maximization. Working paper, Yale University, 2009.

[Seuken *et al.*, 2008] Sven Seuken, Ruggiero Cavallo, and David C. Parkes. Partially-synchronized DEC-MDPs in dynamic mechanism design. In *Proc. 23rd National Conference on Artificial Intelligence (AAAI'08)*, pages 162–169, 2008.

[Vulcano *et al.*, 2002] G. Vulcano, Garret van Ryzin, and C. Maglaras. Optimal dynamic auctions for revenue management. *Management Science*, 48:1388–1407, 2002.

Appendix

Definition 8 (dynamic Groves mechanism)**.** *A dynamic mechanism $M = (\pi_\theta, x_\theta)$ in an environment with a fixed, accessible population and dynamic type is a dynamic Groves mechanism, if:*

(i) policy π_θ is efficient with respect to the reported type profile, and

(ii) each agent i's expected discounted transfer given state profile θ^t (including state profile s^t), strategy σ_i, history h^{t-1}, and given that agents $\neq i$ follow a truthful strategy in period t and forward, is $V_{-i}(h^{t-1}, s^t, \pi_\theta, \sigma_i) - C_i(\theta^t_{-i})$, where $C_i(\theta^t_{-i})$ is independent of agent i's strategy in this period and forward.

Lemma 3. *A dynamic Groves mechanism in an environment with a fixed, accessible population and dynamic type is efficient and wp-EPIC.*

Proof. Let π^*_θ denote the efficient policy given type profile θ. Assume for contradiction that wp-EPIC fails. In particular, we can assume by the one-shot deviation principle (see the proof of Lemma 1 for a discussion) that there is some history h^{t-1}, some type profile θ^t, and some strategy $\sigma_i \neq \sigma^*_i$ that provides more flow payoff than being truthful while deviating only in the current period, where the other agents are truthful in period t and forward. Let $\widehat{\theta}^t_i = \sigma_i(h^{t-1}, \theta^t_i)$ denote this type report, with $\widehat{s}^t_i$ the state report. Then by properties (i) and (ii), we must have

$$V_i(h^{t-1}, s^t, \pi^*_\theta, \sigma_i) + [V_{-i}(s^{t-1}, \theta^t, \pi^*_\theta, \sigma_i) - C_i(\theta^t_{-i})] > \\ V_i(s^t, \pi^*_\theta) + [V_{-i}(s^t, \pi^*_\theta) - C_i(\theta^t_{-i})], \qquad (15)$$

where the terms on the RHS follow from the efficiency of the policy when agent i is truthful. Collecting terms on both sides, we have $V(h^{t-1}, s^t, \pi^*_\theta, \sigma_i) > V(s^t, \pi^*_\theta)$. But now, we can construct policy π', by setting π' equal to π^*_θ in every state profile except for s^t, where we define $\pi'(s^t) = \pi^*_\theta(\widehat{s}^t_i, s^t_{-i})$. With this, we have $V(s^t, \pi') = V(h^{t-1}, s^t, \pi^*_\theta, \sigma_i) > V(s^t, \pi^*_\theta)$, these flow values evaluated with respect to the state dynamics associated with type profile θ, and a contradiction with the efficiency of π^*_θ. ☐

Proof of Theorem 1. Property (i) in Lemma 3 holds by construction. Fix some agent i, period t, type profile θ^t (including state profile s^t), and assume the other agents are truthful. The flow transfer to agent i, given strategy σ_i, is

$$\mathbb{E}_{s^{t..K}}\left[\sum_{k=t}^{K} \gamma^{k-t} \left(r_{-i}(s^k_{-i}, a^k) + \gamma \cdot \mathbb{E}_{s'}[V_{-i}(s', \pi^*_{\theta_{-i}}) \,|\, \widehat{s}^k, a^k] - V_{-i}(\widehat{s}^k, \pi^*_{\theta_{-i}}) \right) \;\Big|\; s^t, \tau, \sigma_i \right], \qquad (16)$$

where $s^k = \tau(s^{k-1}, a^k)$ for $k > t$, $a^k = \pi^*_\theta(\widehat{s}^k)$, $\widehat{s}^k = (\widehat{s}^k_i, s^k_{-i})$, $\widehat{s}^k_i$ is agent i's reported state in period k given strategy σ_i, $s' = \widehat{\tau}^k(\widehat{s}^k, a^k)$, $\widehat{\tau}^k = (\widehat{\tau}^k_i, \tau_{-i})$, and $\widehat{\tau}^k_i$ is agent i's reported transition function in period k.

The second term simplifies as,

$$\gamma \cdot \mathbb{E}_{s'} \left[V_{-i}(s', \pi^*_{\theta_{-i}}) \,\middle|\, s' = \widehat{\tau}^k(\widehat{s}^k, a^k) \right] = \gamma \cdot \mathbb{E}_{s'} \left[V_{-i}(s', \pi^*_{\theta_{-i}}) \,\middle|\, s' = \tau(s^k, a^k) \right], \quad (17)$$

by private values and independent state transitions, conditioned on the action a^k, and can be evaluated on the true distribution of next states given a^k. Similarly, the third term is equivalent to $V_{-i}(s^k, \pi^*_{\theta_{-i}})$ because agent i's reported state does not affect the flow value to other agents.

Returning to (16) and extracting out the sum over the first term, reversing the second and third terms, and simplifying the second term in this way, we have flow transfer,

$$V_{-i}(h^{t-1}, s^t, \pi^*_\theta, \sigma_i) - \mathbb{E}_{s^{t..K}} \left[\sum_{k=t}^{K} \gamma^{k-t} \left(V_{-i}(s^k, \pi^*_{\theta_{-i}}) - \gamma V_{-i}(s^{k+1}, \pi^*_{\theta_{-i}}) \right) \,\middle|\, s^t, \tau, \sigma_i \right], \tag{18}$$

and telescoping out, and noting that $V_{-i}(s^{K+1}, \pi^*_{\theta,-i})$ is necessarily zero, we have flow transfer,

$$V_{-i}(h^{t-1}, s^t, \pi^*_\theta, \sigma_i) - V_{-i}(s^t, \pi^*_{\theta_{-i}}), \tag{19}$$

and this completes the proof, because $V_{-i}(s^t, \pi^*_{\theta_{-i}})$ is independent of agent i's strategy forward from this period. □

Proof of Theorem 2. Property (i) in Lemma 1 holds by construction. In establishing property (ii), consider first a simplified mechanism in which transfer $x_{\theta,i}(\mathsf{s}^t)$ can be made directly, in every period, irrespective of whether or not an agent is accessible. Noting that private values and independent type transitions, conditioned on actions, are retained, and the equivalent form of (7) and (10), it is straightforward (by symbolic substitution) to adopt the proof of Theorem 1 to establish that the flow transfer to agent i is, $V_{-i}(h^{t-1}, \check{\mathsf{s}}^t, \pi^*_\theta, \sigma_i) - V_{-i}(\check{\mathsf{s}}^t, \pi^*_{\theta_{-i}})$, where the second term is independent of agent i's strategy in period t and forward.

We now establish through a simple accounting argument that catch-up payments make the expected discounted transfer to agent i, forward from any period t in which it is accessible, equivalent, for any strategy σ_i, to the expected discounted transfer if it received transfers $x_{\theta,i}(\mathsf{s}^t)$ in every period. For this, consider any sequence of belief types, $\vartheta^t_{-i}, \ldots, \vartheta^K_{-i}$, held by the center about agents $\neq i$ over periods $t, \ldots, K$, under policy π^*_θ, and given that agent i follows strategy σ_i while the other agents are truthful. Let $\theta^t_i, \ldots, \theta^K_i$ denote the corresponding sequence of agent i types. Let $F(\sigma_i) \subseteq \{t, \ldots, K\}$ denote the time periods in which agent i first becomes accessible again, after one or more periods of (reported) inaccessibility. Let $H(\sigma_i) \subseteq \{t, \ldots, K\}$ denote the time periods when agent i reports a non-null type. The realized, total discounted transfer to agent i forward from period t, in the generalized dynamic-VCG mechanism is,

$$\sum_{\substack{k=t\\ k\in H(\sigma_i)\setminus F(\sigma_i)}}^{K} \gamma^{k-t}\cdot x_{\theta,i}(\Gamma(\sigma_i(\theta_i^k)),\mathsf{s}_{-i}^k) \;+\; \sum_{\substack{k=t\\ k\in F(\sigma_i)}}^{K}\gamma^{k-t}\cdot \sum_{k'=k-\delta(k)}^{k}\frac{x_{\theta,i}(\Gamma(\sigma_i(\theta_i^{k'})),\mathsf{s}_{-i}^{k'})}{\gamma^{k-k'}}$$

$$= \sum_{k=t}^{K}\gamma^{k-t}\cdot x_{\theta,i}(\Gamma(\sigma_i(\theta_i^k)),\mathsf{s}_{-i}^k), \tag{20}$$

where $\delta(k) > 0$ is the number of contiguous periods the agent reported itself inaccessible before becoming accessible, $\Gamma(\sigma_i(\theta_i))$ assigns probability 1 to the reported state when $\sigma_i(\theta_i) \neq \phi$ and is defined by Bayes rule otherwise, belief state s_{-i}^k is that associated with belief type ϑ_{-i}^k, and where the dependence of strategy σ_i on history is dropped for notational simplicity. Eq. (20) follows from simple algebra, together with Assumption 1 and observing that $x_{\theta,i}(\mathsf{s}) = 0$ when agent i is not pivotal given belief state s, and so it is immaterial if agent i remains inaccessible in period K because by Assumption 1 its catch-up transfer would be zero in any case. Given that this equivalence holds for any realized sequence of belief types then it also holds in expectation, and this completes the proof. □

Proof of Theorem 3. Interpreting the requirements of Lemma 1 in this environment, we need (i) that the policy π followed by the mechanism is efficient with respect to reported types and the center's model of the arrival process, and (ii) the flow transfer *to an accessible agent* i forward from any type profile θ^t, for any history h^{t-1} and for any strategy σ_i, and given that agents $\neq i$ follow a truthful strategy from this period forward and that the center's model of the arrival process is correct, is $V_{-i}(h^{t-1}, s^t, \pi_\theta, \sigma_i) - C_i(s_{-i}^t)$, where $C_i(s_{-i}^t)$ is independent of agent i's strategy in this period and forward. Property (i) holds by construction. To establish property (ii), consider a modified dynamic-VCG mechanism in which the transfer (7) is collected in *every* period. Given this, it follows by notational substitution into the proof of Theorem 1 that the flow transfer to an agent forward from some period t in which it is accessible is,

$$V_{-i}(h^{t-1}, s^t, \pi_\theta^*, \sigma_i) - V_{-i}(s^t, \pi_{\theta_{-i}}^*), \tag{21}$$

observing that the independent type transitions property continues to hold because of CIA. This flow transfer satisfies property (ii), with term $V_{-i}(s^t, \pi_{\theta_{-i}}^*)$ agent-independent because of the CIA property. Then, the transfer (7) is zero in any period before an agent's reported arrival because the agent's effect on the actions of the efficient policy is only through the arrival process until its reported arrival. In addition, the transfer is zero after an agent's reported departure because it has zero value for actions, and by CIA, has no effect on the belief about future arrival types. This completes the proof. □

Proof of Theorem 4. The decision policy is efficient and satisfies property (i) of Lemma 2 by construction. For property (ii), fix agent i (still to report its type) and consider a period t in which the agent is active, strategy σ_i, history h^{t-1} and partially-truthful type profile $\breve{\theta}^t$ (with associated state profile $\breve{s}^t$). The expected discounted transfer to i, forward from period t, is

$$\mathbb{E}_{s^{t..K}}\left[\gamma^{k(\sigma_i)-t}\left(V(\widehat{s}^{k(\sigma_i)},\pi^*_\theta) - V_{-i}(s^{k(\sigma_i)}_{-i},\pi^*_{\theta_{-i}}) - \sum_{k=k(\sigma_i)}^{K}\gamma^{k-k(\sigma_i)}\cdot\widehat{r}_i(\widehat{s}^k_i,a^k)\right)\Big|\breve{\theta}^t,\sigma_i\right] \quad (22)$$

$$= \mathbb{E}_{s^{t..K}}\Big[\gamma^{k(\sigma_i)-t}\Big(V_{-i}(\widehat{s}^{k(\sigma_i)},\pi^*_\theta) - V_{-i}(s^{k(\sigma_i)}_{-i},\pi^*_{\theta_{-i}})\Big) \mid \breve{\theta}^t,\sigma_i\Big], \quad (23)$$

where $k(\sigma_i)$ is the period in which agent i reports its type given strategy σ_i. The expectation is taken with respect to a sequence of states, $s^k = \breve{\tau}(s^{k-1},a^k)$ for $k > t$, and $s^t = \breve{s}^t$, where $\breve{\tau}$ and $\breve{s}^t$ are the transition function and state associated with $\breve{\theta}^t$. For action $a^k = \pi^*_\theta(\widehat{s}^k)$, this is selected according to π^*_θ, the efficient policy given $\breve{\theta}^t$ and agent i's type as reported under σ_i, and with joint state $\widehat{s}^k = (\widehat{s}^k_i, s^k_{-i})$, where $\widehat{s}^k_i$ is the state reported by agent i in period k. Reward $\widehat{r}_i$ is that reported by agent i. Finally, $\pi^*_{\theta_{-i}}$ is efficient given $\breve{\theta}^t$. By algebra, this flow becomes,

$$= \mathbb{E}_{s^{t..K}}\left[\gamma^{k(\sigma_i)-t}\cdot V_{-i}(\widehat{s}^{k(\sigma_i)},\pi^*_\theta) + \sum_{k=t}^{k(\sigma_i)-1}\gamma^{k-t}\cdot\breve{r}_{-i}(s^k_{-i},a^k) \mid \breve{\theta}^t,\sigma_i\right]$$
$$- \mathbb{E}_{s^{t..K}}\left[\gamma^{k(\sigma_i)-t}\cdot V_{-i}(s^{k(\sigma_i)}_{-i},\pi^*_{\theta_{-i}}) + \sum_{k=t}^{k(\sigma_i)-1}\gamma^{k-t}\cdot\breve{r}_{-i}(s^k_{-i},a^k) \mid \breve{\theta}^t,\sigma_i\right], \quad (24)$$

where $\breve{r}_{-i}$ is the reward profile to agents $\neq i$ associated with $\breve{\theta}^t$, and this simplifies to $V_{-i}(h^{t-1},\breve{s}^t,\pi^*_\theta,\sigma_i) - V_{-i}(\breve{s}^t,\pi^*_{\theta_{-i}})$, where the first term is equal to the first term in (24) by definition, and the second term follows by private values and CIA, which provides that the actions selected by the efficient policy between the current period and period $k(\sigma_i)$ are the same as would be made under $\pi^*_{\theta_{-i}}$. □