

DYNAMIC MECHANISM DESIGN: AN INTRODUCTION

By

Dirk Bergemann and Juuso Välimäki

August 2017

COWLES FOUNDATION DISCUSSION PAPER NO. 2102



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

<http://cowles.yale.edu/>

Dynamic Mechanism Design: An Introduction*

Dirk Bergemann[†] Juuso Välimäki[‡]

August 22, 2017

Abstract

We provide an introduction into the recent developments of dynamic mechanism design with a primary focus on the quasilinear case. First, we describe socially optimal (or efficient) dynamic mechanisms. These mechanisms extend the well known Vickrey-Clark-Groves and D'Aspremont-Gérard-Varet mechanisms to a dynamic environment. Second, we discuss results on revenue optimal mechanism. We cover models of sequential screening and revenue maximizing auctions with dynamically changing bidder types. We also discuss models of information management where the mechanism designer can control (at least partially) the stochastic process governing the agent's types. Third, we consider models with changing populations of agents over time. This allows us to address new issues relating to the properties of payment rules. After discussing related models with risk-averse agents, limited liability, and different performance criteria for the mechanisms, we conclude by discussing a number of open questions and challenges that remain for the theory of dynamic mechanism design.

KEYWORDS: Dynamic Mechanism Design, Sequential Screening, Dynamic Pivot Mechanism, Bandit Auctions, Information Management,

JEL CLASSIFICATION: D44, D82, D83.

*We acknowledge financial support through NSF Grant 1459899 and a Google Faculty Fellowship. We would like to thank Juuso Toikka and Philipp Strack for many helpful conversations. We would like to thank our students in the Advanced Economic Theory Class at Yale University who saw early versions of this survey and Jaehee Song for valuable research assistance.

[†]Department of Economics, Yale University, New Haven, USA, dirk.bergemann@yale.edu

[‡]Aalto University School of Business, Helsinki, Finland, juuso.valimaki@aalto.fi

1 Introduction

In the analysis of economic environments in which information is dispersed amongst agents, the paradigm of mechanism design has been developed to analyze questions of optimal information collection and resource allocation. The aim of these models is to come up with a framework that is sufficiently flexible to treat applications in various fields of economics and precise enough to yield concrete insights and predictions for the models. Over the last decade, the mechanism design approach has been applied to a variety of dynamic settings. In this survey, we overview the basic questions and modeling issues that arise when trying to extend the static paradigm to the dynamic one. We do not aim at maximal generality of the results that we present but we try to bring about the main ideas in the most natural settings where they arise.

By far the best understood setting for mechanism design is the one with independent private values and quasilinear payoffs. Applications of this model cover models of negotiations, auctions, regulation of public utilities, public goods provision, nonlinear pricing, and labor market contracting, to name just a few. In this survey, we concentrate for the most part on this simplest setting. This will in particular imply that we will not include a comprehensive review of the recent work on dynamic taxation and dynamic public finance which has a strong focus on strictly concave rather than quasilinear payoffs. However in Section 5, we shall provide connections between the two classes of payoff environments and comment on the similarities and differences in the analysis and the results.

It is well-known that in dynamic principal-agent models, private information held by the agent leads to optimal long-term contracts that cannot be replicated by a sequence of short-term contracts. We follow the literature on static mechanism design by allowing the principal to commit at the beginning of the game to a contract that covers the entire length of the relationship. In our main model with quasilinear utilities, private information and the desire to control the information rent accruing to the agent is the only reason preventing the principal from ‘selling the project’ to the agent in the initial period.

Since we insist on full commitment power on part of the mechanism designer throughout this survey, we will largely bypass the vast literature on Coasian bargaining that has the lack of commitment at its heart. We will also not discuss in any detail the closely related literature on dynamic games where the players may engage in interactions while their payoff relevant types are subject to stochastic changes.¹ The models that we cover sometimes feature a single allocation, and some-

¹In repeated and dynamic games, the vector of continuation payoffs plays the role of monetary transfers in

times repeated allocations over time. In all our applications, the types of some agents change in a non-trivial manner across periods. For us, this is the distinguishing feature of dynamic mechanism design.

The leading example for this survey is the problem of selling a given supply of goods over time when the demand for the goods is realized over time. Special cases of problems within this class of models include: (i) repeated sales when the buyer is uncertain about her future valuation for the good or service; (ii) leasing a resource to a number of competing bidders that learn over time their private value of the resource and (iii) selling (once and for all) an indivisible (durable) good to a set of potential bidders whose types change over time.

For concreteness, let us describe the first of these settings in more detail and describe some of the economic questions they pose. Leading examples here might be such common situations as memberships to clubs, such as fitness, or long-term contracts, such as mobile phone contracts or equipment service contracts. At any given point in time the buyer knows how much she values the service, but is uncertain about how future valuations for the service may evolve. From the point of view of the service provider, the question then arises how to attract (and sort) the buyers with different current and future valuations for his services. The menu of possible contract presumably has to allow the buyers to express their private current willingness-to-pay as well as their expectation over future willingness-to-pay. A variety of dynamic contracts are empirically documented, for example in gym memberships and mobile phone contracts, as described in DellaVigna and Malmendier (2006) and Grubb and Osborne (2015), respectively. These include (i) flat rates, in which the buyer only pays a fixed fee regardless of her consumption; (ii) two-part-tariffs in which the buyer selects from a menu of fixed fees and variable price per unit of consumption; and (iii) leasing contracts where the length of the lease term is the object of choice for the consumer. In the subsequent survey we will highlight how these and other features of empirically documented contract varieties may arise as solutions to dynamic mechanism design problems.

The first section in this survey develops mechanisms that achieve a *socially efficient* allocation in these and many other models. For dynamic versions of the pivot mechanism, it is important to keep track of the agents' marginal contributions to social welfare over time. In terms of the examples above, the key is to compute the dynamic externalities imposed on other bidders when the good is allocated to a given bidder.

For *budget balanced* mechanisms, the key is to find a way of balancing the payments of different

mechanism design problems.

agents when their information over the distribution of future payments arrives at different times. Balancing the budget in bilateral trading problems is more subtle than in the static case if the parties learn their types at different points in the allocation game.

For the case of *revenue maximizing* mechanisms, one might guess that the part of private information held by the bidders at the moment of contracting is the only source of information rent: the rest of the stochastic process is uncertain to both the seller and the buyer, and after the initial report the two parties share a common probability distribution on future types. This intuition has been formalized by describing the additional information through a process of orthogonal information and this insight forms a key part of the analysis of revenue maximizing dynamic mechanisms.

The key implications for the revenue maximizing allocation stem from this intuition. For most stochastic processes (e.g. ergodic and strongly mixing processes), knowing the value of the process in period t tells little about the value of the process in $t+k$ for k large. Hence one might conjecture that the private information θ_0 held by the agent at the moment of signing the contract provides little private information on the valuation θ_t for large t . Hence the reasons to distort the allocation for large t in order to extract information rent from the agent at $t = 0$ vanish as t becomes large. Along the revenue maximizing sales mechanism, allocations and trades converge to efficient allocation over time. The key to the analysis lies in finding an appropriate dynamic envelope formula to express the necessary conditions of incentive compatibility for the agent. These conditions result in dynamic virtual surplus formulas where the static Myersonian virtual surplus is modified by an *impulse response* that measures the impact of first period private information on the future types.

With risk aversion and with limited liability, we add different reasons beyond private information for continued contracting between the parties. Much less is known about the optimal solutions in these cases. We outline the dynamic model under these alternative assumptions and discuss some observations that follow from the Myersonian model. Incentives for local deviations, i.e. incentives for reporting private information near the true type in our dynamic direct mechanism can still be computed using a dynamic envelope theorem. The step of solving the relaxed program that reflects these necessary incentive compatibility conditions is a lot more involved due to the lack of quasilinearity. Unfortunately, the step of verifying full incentive compatibility is also a lot harder.

We divide the survey into three parts. The first part covers efficient dynamic mechanisms. This is the dynamic counterpart of the Vickrey-Clark-Groves (VCG) mechanisms and budget balanced d'Aspremont-Gérard-Varet (AGV) mechanisms. We view this setting as the natural first benchmark case for dynamic allocation models in the same way as the VCG mechanisms provide a

natural benchmark for static problems. Examples of potential applications include efficient dynamic auctions, efficient dynamic bilateral trade or bargaining, and dynamic public goods problems.

The second part deals with optimal dynamic mechanisms. The corresponding static models here are the optimal screening models in a principal-agent setting and different models of selling procedures such as auctions. Here we cover sequential screening and revenue maximizing auctions with dynamically changing bidder types. We also discuss nonlinear pricing over time when the type of the buyer changes stochastically. We also comment on the model of information management where the mechanism designer can control (at least partially) the stochastic process governing the agent's types.

The third part considers models with changing populations of agents over time. Obviously this part has no counterpart on the static side. It allows us to ask new questions relating to the properties of the payment rules. For example with changing populations, it makes sense to require that agents receive or make transfers only in the periods when they are alive. These restrictions lead to interesting new findings about the models where efficient outcomes can be achieved. We discuss also some work on dynamic revenue management in this setting.

In the last substantive section of this survey, we consider briefly related models from public finance and financial economics. The key departure in these models is the lack of quasi-linearity. The models in dynamic public finance have concerns of consumption smoothing over risky outcomes at their heart. Hence the models feature agents with strictly concave utilities in consumption and leisure. In addition to the possibility of having risk-averse decision makers, the models in financial economics often feature a limited liability constraint on the transfer rules: owners can pay the managers but managers cannot be asked to make (arbitrarily large) payments to the owners. We discuss the similarities in the analysis and contrast the results of these models with the models under quasilinear utility. Finally, we make some connections to the rapidly growing literature on mechanism design in computer science. Rather than concentrating on the properties of the optimal mechanism for a fixed stochastic model, this literature takes as its objective the task of finding mechanisms that guarantee a good payoff across a variety of different stochastic models.

While the target reader of this survey is an applied economist interested in learning about the applicability of the mechanism design approach to dynamic incentive problems, we develop the analytical arguments to some extent here and also comment on some theoretically subtle issues. The interested reader will find complementary material and more technical detail in the recent textbooks by Börgers (2015) and Gershkov and Moldovanu (2014), and earlier surveys focusing on dynamic

auctions by Bergemann and Said (2011), and more recently by Pavan (2017). Bergemann and Pavan (2015) provide an introduction into recent research in dynamic mechanism design collected in a symposium issue of the *Journal of Economic Theory*.

2 Efficient Dynamic Mechanisms

In this section, our aim is to extend the familiar Vickrey-Clark-Groves (VCG) mechanisms of static mechanism design to a dynamic model. In line with the static model, we insist on private values, independent types and quasilinear utilities. With these assumptions, efficient (in the sense of maximizing the gross social surplus) dynamic VCG mechanisms exist and they are the only ex post incentive compatible efficient mechanisms. With additional assumptions on the type sets, we can demonstrate revenue equivalence results similar to those in the static model.

We start by setting up the general dynamic allocation problem that will be used throughout this survey. We derive the dynamic counterpart to the static pivot mechanism and explain how the dynamic mechanism succeeds in allocating each agent in the model her marginal contribution to social welfare. We discuss then the dynamic version of the budget balanced dynamic d'Aspremont-Gérard-Varet (AGV) mechanism of Athey and Segal (2013). Finally, we discuss extensions to correlated and interdependent values respectively.

2.1 The Dynamic Allocation Problem

We consider an environment with private and independent values in a discrete-time, infinite-horizon model. Agent $i \in \{1, 2, \dots, I\}$ gets flow utility in period $t \in \mathbb{N}$ that depends on the current physical allocation $x_t \in X_t$, the current monetary transfer $p_{i,t} \in \mathbb{R}$, and the type $\theta_{i,t} \in \Theta_i$. The Bernoulli utility function u_i of agent i is quasi-linear in the monetary transfer:

$$u_i(x_t, p_{i,t}, \theta_{i,t}) \triangleq v_i(x_t, \theta_{i,t}) - p_{i,t}.$$

We assume that the type $\theta_{i,t}$ of agent i follows a controlled Markov process on the state space Θ_i . The type vector in period t is given by $\theta_t = (\theta_{1,t}, \dots, \theta_{I,t}) \in \Theta$, with $\Theta = \times_{i=1}^I \Theta_i$. The set of feasible allocations in t may depend on the vector of past allocations $x^t := (x_0, \dots, x_t)$ as in a dynamic auction of a single indivisible object.²

²In order to keep the notation manageable, we do not index X_t by x^t .

There is a common prior $F_i(\theta_{i,0})$ regarding the initial type $\theta_{i,0}$ of each agent i . The current type $\theta_{i,t}$ and the current action x_t define a probability distribution for the state variables $\theta_{i,t+1}$ on Θ_i in the next period. We assume that this distribution can be represented by a transition function (or stochastic kernel)

$$F_i(\theta_{i,t+1} | \theta_{i,t}, x_t).$$

The utility functions $u_i(\cdot)$ and the transition functions F_i are all common knowledge at $t = 0$. The common prior $F_i(\theta_{i,0})$ and the functions $F_i(\theta_{i,t+1} | \theta_{i,t}, x_t)$ are assumed to be independent across agents. At the beginning of each period t , each agent i observes $\theta_{i,t}$ privately. At the end of each period, an action $x_t \in X$ is chosen and payoffs for period t are realized. The asymmetric information is therefore generated by the private observation of $\theta_{i,t}$ in each period t . We observe that by the independence of the priors and the transition functions across i , the information of agent i , $\theta_{i,t+1}$, does not depend on $\theta_{j,t}$ for $j \neq i$. We assume that

$$|v_i(x, \theta_i)| < K,$$

for some $K < \infty$ for all i, x and θ_i .

Since we do not want to introduce gains from trade through differences in intertemporal marginal rates of substitution between the players, we assume that all agents discount the future with a common discount factor δ , $0 < \delta < 1$. When discussing welfare, we adopt the point of view that the mechanism designer has in period t a flow payoff given by

$$u_0(x_t, p_t, \theta_t) = v_0(x_t) + \sum_{i=1}^I p_{i,t}.$$

The socially efficient policy is obtained by maximizing the expected discounted sum of valuations. Notice that if the mechanism designer is not an actual player in the game, this definition of efficiency does not take into account budget deficits and surpluses.

Given the Markovian structure of the type processes, the socially optimal program starting in period t at type vector θ_t can be written as:

$$\begin{aligned} W(\theta_t) &\triangleq \max_{\{x_s\}_{s=t}^{\infty}} \mathbb{E} \left\{ \sum_{s=t}^{\infty} \delta^{s-t} \sum_{i=0}^I v_i(x_s, \theta_{i,s}) \right\} \\ &\triangleq \max_{\{x_s\}_{s=t}^{\infty}} \mathbb{E} \left\{ \sum_{s=t}^{\infty} \delta^{s-t} w(x_s, \theta_s) \right\}, \end{aligned}$$

where we have defined

$$w(x_s, \theta_s) \triangleq \sum_{i=0}^I v_i(x_s, \theta_{i,s}).$$

Notice that the dynamic dependence across the periods is now buried in the expectations operator. Since the distribution of $\theta_{i,t+1}$ depends on x_t , the expectation over future types depends on the current choice of x .

Alternatively, we can represent the social program in its recursive form:

$$\begin{aligned} W(\theta_t) &= \max_{x_t} \mathbb{E} \{ w(x_t, \theta_t) + \delta \mathbb{E} W(\theta_{t+1}) \}. \\ \text{s.t. } \theta_{t+1} &\sim F(\cdot | \theta_t, x_t) \end{aligned}$$

The socially efficient policy is denoted by $x^* = \{x_t^*\}_{t=0}^\infty$. The social externality cost of agent i is determined by the social value in the absence of agent i :

$$W_{-i}(\theta_t) \triangleq \max_{\{x_s\}_{s=t}^\infty} \mathbb{E} \left\{ \sum_{s=t}^\infty \delta^{s-t} \sum_{j \neq i} v_j(x_s, \theta_{j,s}) \right\} \triangleq \max_{\{x_s\}_{s=t}^\infty} \mathbb{E} \left\{ \sum_{s=t}^\infty \delta^{s-t} w_{-i,s}(x_s, \theta_s) \right\},$$

where the efficient policy when agent i is excluded is denoted by $x_{-i}^* = \{x_{-i,t}^*\}_{t=0}^\infty$, and $w_{-i,t}(x_t, \theta_t)$ denotes the flow social welfare to the society where i has been excluded. The *marginal contribution* $M_i(\theta_t)$ of agent i at signal θ_t is defined by:

$$M_i(\theta_t) \triangleq W(\theta_t) - W_{-i}(\theta_t). \tag{1}$$

The marginal contribution of agent i is the change in the social value due to the addition of agent i . To fix ideas, we give a canonical example of a dynamic auction with uncertainty about future values.

Example: Option Value and Irrevocable Allocation We consider the assignment of a single indivisible object between two bidders $i \in \{1, 2\}$ at a unique, and then irrevocable, time period t . Thus the object can only be assigned once and only has value in a single period. The question though is to whom should the object be allocated and in which time period. The allocation x_t in period t is a vector $(x_{1,t}, x_{2,t})$ with $x_{i,t} \in \{0, 1\}$, where $x_{i,t} = 1$ denotes allocating the object to i in period t . The set of feasible allocations in period t is $X_t = \{(0, 1), (1, 0), (0, 0)\}$ if $x_s = (0, 0)$ for all $s < t$. Otherwise $x_t = \{(0, 0)\}$.

Bidder 1 has a constant value $\theta_{1,t} = \theta_1 \in [0, 1]$ for the object for all t . By contrast, bidder 2 learns her value in period 1, and $\theta_{2,t} \in \{0, 1\}$ for $t > 0$. Thus, there might be value of postponing

the assignment decision until period 1, rather than assigning and consuming the object in period 0. The information that arrives in period 1 thus suggest that there might be an option value from postponing the assignment until period 1. The private information of bidder 2 in period $t = 0$ is the probability of having value 1 from period 1 onwards: $\theta_{2,0} = \Pr\{\theta_{2,t} = 1 \text{ for } t > 0 | \theta_{2,0}\}$.

The payoff to bidder i from a feasible allocation is $\sum_t \delta^t x_{i,t} \theta_{i,t}$. In words, bidder i gets her (expected) payoff of $\theta_{i,t}$ if she is allocated the good in t . We see immediately that

$$W_{-1}(\theta_1, \theta_{2,t}) = \Pr\{\theta_{2,t} = 1 \text{ for } t > 0 | \theta_{2,t}\} = \theta_{2,t}$$

and

$$W_{-2}(\theta_1, \theta_{2,t}) = \theta_1.$$

To compute the efficient, policy, we compare the payoff from immediate allocation to waiting until $t = 1$ before allocating. Waiting yields social welfare

$$\delta \mathbb{E} \max\{\theta_1, \theta_{2,0}\} = \delta(\theta_{2,0} + (1 - \theta_{2,0}) \theta_1).$$

Immediate allocation yields $\max\{\theta_1, \theta_{2,0}\}$. Hence the optimal decision is to allocate immediately to 1 if

$$\theta_1 \geq \frac{\delta \theta_{2,0}}{1 - \delta + \delta \theta_{2,0}},$$

and to allocate immediately to 2 if

$$\theta_{2,0} \geq \frac{\delta \theta_1}{1 - \delta + \delta \theta_1}.$$

If $x_{i,0}^* = 1$, then $M_i(\theta_0) = \theta_{i,0} - \theta_{-i,0}$ and $M_{-i}(\theta_0) = 0$. If $x_{1,0}^* = x_{2,0}^* = 0$, then $M_1(\theta_0) = \delta(\theta_{2,0} + (1 - \theta_{2,0}) \theta_1) - \theta_{2,0}$ and $M_2(\theta_0) = \delta(\theta_{2,0} + (1 - \theta_{2,0}) \theta_1) - \theta_1$. In this case, we can also compute $M_1(\theta_1) = \max\{\theta_1, \theta_{2,1}\} - \theta_1$, and $M_2(\theta_1) = \max\{\theta_1, \theta_{2,1}\} - \theta_{2,1}$. The marginal contributions beginning with period 1 coincide with the static marginal contributions.

2.1.1 Mechanism and Equilibrium

We focus attention on direct mechanisms which implement the socially efficient policy x^* . In other words, we seek a transfer rule that induces all the agents to report their type truthfully when the physical allocation is given by the efficient allocation rule given the reported types.

In the static case, the task is transparent. By the taxation principle, we can view the reporting problem as an equivalent problem of letting each agent choose her most preferred allocation given the reports of other agents. By promising each agent the entire social surplus from the efficient

choice given the reported types of other agents and her own true type, she has always the right incentives to choose the efficient allocation. In other words, she has a dominant strategy to report her type truthfully in the direct mechanism. We shall see that some of the nice properties of the static VCG mechanisms fail. Dynamic efficient mechanisms will not in general be in dominant strategies. Furthermore they are not detail free in the sense that details of the transition probabilities $F_i(\theta_{i,t} | \theta_{i,t-1}, x_t)$ matter for the efficient choice of allocation in each t .

In a dynamic direct mechanism every agent i is asked to report her type $\theta_{i,t}$ in every period t . We say that the dynamic direct mechanism is truthful if the reported type $r_{i,t} \in \Theta_i$ coincides with the true type for all $i, t, \theta_{i,t}$. The dynamic revelation principle as stated in Myerson (1986) shows that there is no loss of generality in restricting attention to direct mechanisms where the agents report their information truthfully on the equilibrium path.

The mechanism designer chooses how much of the information in the reports and allocations to disclose to the players. If $x_t = (x_{1,t}, \dots, x_{I,t})$ and $v_i(x_t, \theta_{i,t}) = v_i(x_{i,t}, \theta_{i,t})$, for all i, t , then the minimal information to disclose to player i would be $x_{i,t}$. It is clear that restricting the information available to agent i makes it easier to satisfy the incentive compatibility constraints for that player. In the current section, we can find efficient mechanisms that satisfy very strong notions of incentive compatibility and as a result, we assume that the entire vector of reports and allocations is disclosed at each t .

With this assumption in place, the public history in period t is a sequence of reports and allocations until period $t - 1$, or $h_t = (r_0, x_0, r_1, x_1, \dots, r_{t-1}, x_{t-1})$, where each $r_s = (r_{1,s}, \dots, r_{I,s})$ is a report profile of the I agents. The set of possible public histories in period t is denoted by H_t . The sequence of reports by the agents is part of the public history and we assume that the past reports of each agent are observable to all the agents. The private history of agent i in period t consists of the public history and the sequence of private observations until period t , or $h_{i,t} = (\theta_{i,0}, r_0, x_0, \theta_{i,1}, r_1, x_1, \dots, \theta_{i,t-1}, r_{t-1}, x_{t-1}, \theta_{i,t})$. The set of possible private histories in period t is denoted by $H_{i,t}$. An (efficient) dynamic direct mechanism is given by a family of allocations and monetary transfers, $\{x_t^*, p_t\}_{t=0}^\infty$: $x_t^* : \Theta \rightarrow \Delta(X_t)$, and $p_t : H_t \times \Theta \rightarrow \mathbb{R}^I$. With the focus on efficient mechanisms, the allocation x_t^* depends only on the current (reported) type $r_t \in \Theta$. In contrast, the transfer p_t may depend on the entire history of reports and actions.

A (pure) reporting strategy for agent i in period t is a mapping from the private history into the type space: $r_{i,t} : H_{i,t} \rightarrow \Theta_i$. For a given mechanism, the expected payoff of agent i from reporting

strategy $r_i = \{r_{i,t}\}_{t=0}^\infty$ given the strategies $r_{-i} = \{r_{-i,t}\}_{t=0}^\infty$ is:

$$\mathbb{E} \sum_{t=0}^{\infty} \delta^t [v_i(x^*(r_t), \theta_{i,t}) - p_i(h_t, r_t)].$$

Given the mechanism $\{x_t^*, p_t\}_{t=0}^\infty$ and the reporting strategies r_{-i} , the optimal strategy of bidder i can be stated recursively:

$$V_i(h_{i,t}) = \max_{r_{i,t} \in \Theta_i} \mathbb{E} \{v_i(x_t^*(r_{i,t}, r_{-i,t}), \theta_{i,t}) - p_i(h_t, r_{i,t}, r_{-i,t}) + \delta V_i(h_{i,t+1})\}.$$

The value function $V_i(h_{i,t})$ expresses the continuation value of agent i given the current private history $h_{i,t}$. We say that a dynamic direct mechanism is *interim incentive compatible*, if for every agent and every history, truthtelling is a best response given that all other agents report truthfully. We say that the dynamic direct mechanism is *periodic ex-post incentive compatible* if truthtelling is a best response regardless of the history and the current type of the other agents.

In the dynamic context, the notion of ex-post incentive compatibility is qualified by *periodic* as it is ex-post with respect to all signals received until and including period t , but not ex-post with respect to signals arriving *after* period t . The periodic qualification arises in the dynamic environment as agent i may receive information at some later time $s > t$ such that in retrospect she would wish to change the allocation choice in t and hence her report in t .

2.1.2 Types, Allocations and Time Horizon

We have not said much about the interpretation of the types $\theta_{i,t}$. Some authors separate two classes of models: models with exogenous and endogenous types depending on whether the allocation decisions have an impact on the distribution of future types. An example of exogenous types could be a model of procuring goods from firms subject to autocorrelated privately observed cost shocks. In this case, we could take $\theta_{i,t} \in \mathbb{R}$, with

$$\theta_{i,t+1} = \gamma\theta_{i,t} + \varepsilon_{i,t+1},$$

where the $\varepsilon_{i,t}$ are i.i.d. shocks.

An example of the second class of models, consider an employer who learns privately about the (firm-specific) productivity ω_i of a worker in periods when she employs the worker. In this case it would be natural to take $\theta_{i,t}$ to be the probability distribution of the true productivity ω_i . The worker produces a output $y_{i,t}$ that is privately observed by the employee and distributed according

to a c.d.f. $G(\cdot | \omega_i)$. The transition for the types, $\theta_{i,t+1}$ is obtained from $\theta_{i,t}$ and $y_{i,t}$ by Bayes' rule. For example in a normal learning model, the prior on ω_i would be a normal random variable and each $y_{i,t}$ is normally distributed with mean ω_i and a known variance σ^2 . For the periods where the worker is not employed, we have simply $\theta_{i,t+1} = \theta_{i,t}$.

At the cost of some notational inconvenience, we could have allowed the payoffs and the transitions to depend on the full history of allocations: $x^t = (x_0, \dots, x_t)$. It will become clear that none of the results would change as a result of this more general formulation. Hence we can easily accommodate models of learning by doing, change for variety etc.

We can also accommodate the finite-horizon version of the model. This entails simply specifying a certain transition at some T to an absorbing type vector θ_T with the understanding that $v_i(x, \theta_{i,T+s}) = 0$ for all i , all $s > 0$ and all $x \in X$.

For the remainder of this section, we concentrate on particular efficient mechanisms that have further desirable properties. In the next subsection, we describe the dynamic pivot mechanism, introduced in Bergemann and Välimäki (2010), that ensures that each agent's payoff in the mechanism corresponds to her marginal contribution to the societal welfare as defined above after all histories. In the dynamic pivot mechanism, all agents have the correct societal incentives to engage in private investments in e.g. increasing their own payoffs through cost reducing investments. We also consider the dynamic counterpart of the AGV mechanism where the focus shifts towards budget balance. For dynamic bargaining processes and dynamic problems of public goods provision, these considerations are of obvious importance just as they are in the static case.

2.2 The Dynamic Pivot Mechanism

We now construct the dynamic pivot mechanism for the general model described in Subsection 2.1. The marginal contribution of agent i is her contribution to the social value. In the dynamic pivot mechanism, we show that the marginal contribution will also be equal to the equilibrium payoff that agent i can secure for herself along the socially efficient allocation. If agent i receives her marginal contribution in every continuation game of the mechanism, then she should receive the *flow* marginal contribution $m_i(\theta_t)$ in each period. The flow marginal contribution accrues incrementally over time and is defined recursively:

$$M_i(\theta_t) = m_i(\theta_t) + \delta \mathbb{E} M_i(\theta_{t+1}).$$

The flow marginal contribution can be expressed directly in terms of the social value functions,

using the definition of the marginal contribution given in (1), as:

$$m_i(\theta_t) \triangleq W(\theta_t) - W_{-i}(\theta_t) - \delta \mathbb{E}[W(\theta_{t+1}) - W_{-i}(\theta_{t+1})]. \quad (2)$$

The continuation payoffs of the social programs with and without i , respectively, may be governed by different transition probabilities as the respective social decisions in period t , $x_t^* \triangleq x^*(\theta_t)$ and $x_{-i,t}^* \triangleq x_{-i}^*(\theta_{-i,t})$, may differ. The expected continuation value of the socially optimal program, conditional on current allocation x_t and current state θ_t is:

$$W(\theta_{t+1} | x_t, \theta_t) \triangleq \mathbb{E}_{F(\theta_{t+1}; x_t, \theta_t)} W(\theta_{t+1}),$$

where the transition from state θ_t to state θ_{t+1} is controlled by the allocation x_t . For notational ease we omit the expectations operator \mathbb{E} from the conditional expectation. We adopt the same notation for the marginal contributions $M_i(\cdot)$ and the individual value functions $V_i(\cdot)$. The flow marginal contribution $m_i(\theta_t)$ is expressed as:

$$m_i(\theta_t) = w(x_t^*, \theta_t) - w_{-i}(x_{-i,t}^*, \theta_t) + \delta [W_{-i}(\theta_{t+1} | x_t^*, \theta_t) - W_{-i}(\theta_{t+1} | x_{-i,t}^*, \theta_t)].$$

A monetary transfer $p_i^*(\theta_t)$ such that the resulting flow net utility matches the flow marginal contribution leads agent i to internalize her social externalities:

$$p_i^*(\theta_t) \triangleq v_i(x_t^*, \theta_{i,t}) - m_i(\theta_t). \quad (3)$$

We refer to $p_i^*(\theta_t)$ as the transfer of the dynamic pivot mechanism. The transfer $p_i^*(\theta_t)$ depends only on the current report θ_t and not on the entire public history h_t . We can express $p_i^*(\theta_t)$ in terms of the flow utilities and the social continuation values:

$$p_i^*(\theta_t) = w_{-i}(x_{-i,t}^*, \theta_t) - w_{-i}(x_t^*, \theta_t) + \delta [W_{-i}(\theta_{t+1} | x_{-i,t}^*, \theta_t) - W_{-i}(\theta_{t+1} | x_t^*, \theta_t)].$$

Notice that in contrast to the static transfer payment, the reported type of agent i has also an indirect effect through $\delta W_{-i}(\theta_{t+1} | x_t^*, \theta_t)$. This reflects the intertemporal internalization of future externalities that is necessary for aligning the incentives with the planner's dynamic optimum. Given that we started our construction from the requirement that each agent receives her full marginal contribution $W(\theta_t) - W_{-i}(\theta_t)$, we are obviously in the realm of (dynamic) VCG mechanisms.

Theorem 1 (Dynamic Pivot Mechanism)

The dynamic pivot mechanism $\{x_t^, p_t^*\}_{t=0}^\infty$ is ex-post incentive compatible and individually rational.*

In Bergemann and Välimäki (2010), we give conditions for the uniqueness of the above payment rule. The dynamic pivot mechanism has properties that other VCG schemes do not necessarily have. All payments are online in the sense that once an agent is irrelevant for future allocations, she is not asked to make any payments. Furthermore the property of equating equilibrium payoffs with marginal contributions gives the individual agents the socially correct incentives to engage in privately costly investments in θ_i . For a class of dynamic auctions, Mierendorff (2013) develops a dynamic Vickrey auction that satisfies a strong ex post individual rationality requirement.

We continue with the example from the previous subsection to illustrate how the payments in the dynamic pivot mechanism are computed.

Example (Continued) For $x_{i,0}^* = 1$, we have $m_i(\theta_0) = \theta_{i,0} - \theta_{-i,0}$, and $p_i^*(\theta_0) = \theta_{-i,0}$, $p_{-i,0}(\theta_0) = 0$. For $x_0^* = (0, 0)$ we have $m_i(\theta_0) = -(1 - \delta)\theta_{-i,0}$ and $p_i(\theta_0) = (1 - \delta)\theta_{-i,0}$.

Hence we see that the dynamic pivot mechanism asks both bidders to make a positive payment in period 0 if the good is not sold immediately. This reflects the delay externality that their inclusion in the model imposes on the other party.

Since we have assumed independent types, additional assumptions on the connectedness of the type spaces and payoff functions guarantee a dynamic revenue equivalence result. By imposing an individual rationality or participation constraint for the agents, it is often possible to show as in the static setting that the dynamic pivot mechanism results in the maximal monetary surplus among all efficient mechanisms. A negative surplus in the dynamic pivot mechanism then implies an impossibility result mirroring the static Myerson-Satterthwaite theorem on budget balanced efficient dynamic mechanisms that satisfy incentive compatibility and individually rationality.

2.3 An Efficient Dynamic Mechanism

One problem with the dynamic pivot mechanism is that just like its static counterpart, it does not satisfy budget balance. To remedy this, Athey and Segal (2013) develop a dynamic version of the AGV mechanism. The starting point for this construction is the dynamic team mechanism where each agent's individual payoff is augmented by a transfer that makes her total payoff equal to the maximal social surplus in the model. This part of the construction follows along the familiar lines of the static setting where the type vector θ is drawn for a single period. The static VCG mechanisms take the form $(x^*(\theta), p(\theta))$, where for all i and all θ , we have

$$p_i(r) = -w_{-i}(x^*(r_i, r_{-i}), r_{-i}) + \pi_i(r_{-i}),$$

where the second component of the transfer function, $\pi_i(r_{-i})$ is a term that for agent i only depends on the reports of the other agents, r_{-i} . In other words, the transfer of agent i depends on her own announcement only through its impact on the other players' payoffs via the efficient allocation rule. By a simple application of the one-shot deviation principle, one sees that by setting

$$p_{i,t}(r_{i,t}, r_{-i,t}) = -w_{-i}(x^*(r_{i,t}, r_{-i,t}), r_{-i,t}),$$

each agent responds optimally by announcing her type truthfully. This is not surprising: if future announcements are truthful, each agent internalizes the social planners payoffs at all stages in the game. Notice that since all players are receiving the entire social surplus here, this mechanism provides the right answers for truthful reports even in the case of correlated types.

Obviously this mechanism does not satisfy budget balance. The AGV mechanism in the static setting is constructed using the idea that by taking expectations over r_{-i} in the expression $w_{-i}(x^*(r_i, r_{-i}), r_{-i})$, we can take (with the understanding that $I + 1 = 1$):

$$\pi_{i+1}(r_{-(i+1)}) = \mathbb{E}_{r_{-i}} w_{-i}(x^*(r_i, r_{-i}), r_{-i}).$$

With this specification, the budget is clearly balanced and incentive compatibility holds in the sense of Bayesian incentive compatibility.

For the dynamic mechanism, the construction is not quite as simple. Supposing that the incentive payments are made as above based on the expectations over other players' types, but the realizations of the other types become available before one's own announcement, the simple AGV -mechanism is no longer incentive compatible. In order to secure incentive compatibility, Athey and Segal (2013) proceed as follows. Define

$$\gamma^i(\theta_{i,t+1}, \theta_t) = - \left\{ \mathbb{E}_{\tilde{\theta}_{-i,t+1}} \left[W_{-i}(\theta_{i,t+1}, \tilde{\theta}_{-i,t+1} | \theta_t, \mathbf{x}^*) \right] - \mathbb{E}_{\tilde{\theta}_{t+1}} \left[W_{-i}(\tilde{\theta}_{t+1} | \theta_t, \mathbf{x}^*) \right] \right\}.$$

By specifying

$$p_{i,t}(r_{i,t}, r_{-i,t}) = \gamma^i(r_{i,t+1}, r_t) - \gamma^{i+1}(r_{i,t+1}, r_t),$$

the mechanism is obviously budget balanced.

The incentive payment $\gamma^i(\theta_{i,t+1}, \theta_t)$ now reflects the expected change in W_{-i} (conditional on past announcements r_t) that arises from the announcement $\theta_{i,t+1}$. This cannot be manipulated by players different from i in equilibrium. It should be noted that these mechanisms will not satisfy individual rationality in general (since otherwise they would yield a contradiction to the Myerson-Satterthwaite theorem). Dynamic AGV-like mechanisms have been used in e.g. the analysis of

optimal allocation of sharing of the capacity in a joint project when information regarding future profits arrives over time in Kurikbo, Lewis, Liu, and Song (2017).³

2.4 Interdependent Values and Correlation

Dynamic VCG mechanisms have been generalized to cover the case of correlated and interdependent values in Liu (2013). Correlation across agents allows for the use of dynamic versions of mechanisms in the style of Cremer and McLean (1985), (1988). In a dynamic setting, it is possible to use the intertemporal correlation of the reports of the agents and this allows for new types of implementations of the efficient allocation path.⁴ The paper also covers the case of interdependent but independent values. For that case, it is well-known that efficient mechanisms fail to exist in the static setting (see e.g. Dasgupta and Maskin (2000), Jehiel and Moldovanu (2001)). For the case of single dimensional types and appropriate single crossing conditions for the agents' payoffs, Liu (2013) develops a dynamic version of the generalized VCG mechanism.⁵

3 Optimal Dynamic Mechanisms

The analysis of the revenue-maximizing contract in an environment where the agent's private information may change over time appears first in Baron and Besanko (1984). They consider a two-period model of a regulator facing a monopolist with unknown, but in every period, constant marginal cost. Besanko (1985) offers an extension to a finite horizon environment with a general cost function, where the unknown parameter is either i.i.d. over time or follows a first-order autoregressive process. Since these early contributions, the literature has developed considerably in recent years. Courty and Li (2000) consider the revenue-maximizing contract in a sequential screening problem where the preferences of the buyer change over time but the allocation decision takes place

³In their particular problem, Kurikbo, Lewis, Liu, and Song (2017) find a version of the mechanism that can also handle individual rationality constraints and incentives incentives for efficient investments that affect other agents' payoffs.

⁴The idea that we can strengthen the incentive constraints of current announcements by using future realizations of correlated signals is reminiscent of the use of multiple signals in Deb and Mishra (2014) and Mezzetti (2004), (2007).

⁵See Bergemann and Välimäki (2002) for the definition of the static generalized VCG mechanism for allocation problems with single-dimensional types.

only at a single point in time.⁶ Battaglini (2005) considers a quantity discriminating monopolist who provides a menu of choices in each period to a consumer whose valuation can change over time according to a commonly known Markov process with two states. In contrast to the earlier work, he explicitly considers an infinite time horizon and shows that the distortion due to the initial private information vanishes over time.

Pavan, Segal, and Toikka (2014) consider a general environment in an infinite horizon setting and allowing for general allocation problems, encompassing the earlier literature (with continuous type spaces). They obtain necessary conditions for incentive compatibility and present a variety of sufficient conditions for revenue-maximizing contracts for specific classes of environments. They also observe the beneficial implications of time separable environments for a tighter characterization of the optimal contract. Eső and Szentes (2017) and Li and Shi (2017) discuss the amount of information rent that the optimal contract leaves to the agent. In particular, Eső and Szentes (2017) show that under suitable restrictions on the implementability of the optimal allocation, the agent gains no additional information rent from the new information that she gets after signing the contract.

We start by recalling some notions from static optimal mechanisms. Our first dynamic model deals with the simplest model of sequential screening as developed in Courty and Li (2000). We review this material rather quickly since Chapter 11 by Krämer and Strausz in Börgers (2015) contains an excellent textbook overview of the material. Nevertheless, this material serves as a useful preview of the issues that arise also in more complicated models. In order to make any progress, we have to get some characterizations for incentive compatibility in the dynamic setting.

3.1 Preliminaries from Static Mechanism Design

Not surprisingly, we need to rely on tools that make the static mechanism design problem tractable: the revenue equivalence theorem and single crossing properties of the objective function. Hence we follow in the tradition that begun with the early contribution to the dynamic regulation problem in Baron and Besanko (1984), Besanko (1985) and Riordan and Sappington (1987). In the static principal-agent setting, a direct mechanism $(x(\theta), p(\theta))$ is incentive compatible if for all types $\theta \in \Theta$ and all reports $r \in \Theta$, we have

$$U(\theta; \theta) \triangleq v(x(\theta), \theta) - p(\theta) \geq v(x(r), \theta) - p(r) \triangleq U(\theta; r),$$

⁶Eső and Szentes (2007) rephrased the two period sequential screening problem by showing that the additional signal arriving in period two can always be represented by a signal that is orthogonal to the signal in period one.

and let

$$V(\theta) \triangleq U(\theta; \theta).$$

The envelope theorem by Milgrom and Segal (2002) gives the following necessary condition for incentive compatibility in the case where the set of possible types is an interval of the real line: $\Theta = [\underline{\theta}, \bar{\theta}]$.

Theorem 2 (Revenue Equivalence Theorem)

Assume that $v(x, \cdot)$ is differentiable for all $x \in X$ and that there exists a $K < \infty$ such that for all $x \in X$ and all θ ,

$$|v_\theta(x, \theta)| \leq K.$$

Then $V(\theta)$ is absolutely continuous, $V'(\theta) = v_\theta(x(\theta), \theta)$ for almost every θ , and therefore

$$V(\theta) = V(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} v_\theta(x(s), s) ds. \tag{4}$$

This result is called revenue equivalence theorem because we can now pin down the transfers by just determining the physical allocation $x(\theta)$ and the additive constant $V(\underline{\theta})$:

$$p(\theta) = v(x(\theta), \theta) - V(\underline{\theta}) - \int_{\underline{\theta}}^{\theta} v_\theta(x(s), s) ds. \tag{5}$$

With the help of this necessary condition for implementability, we can rewrite the full incentive compatibility requirement as the following *integral monotonicity* condition:

$$\int_r^{\theta} (v_\theta(x(s), s) - v_\theta(x(r), s)) ds \geq 0. \tag{6}$$

Notice that this characterization of implementable allocations remains essentially unchanged even if we allow for multidimensional allocations (these could correspond to the physical allocations of a single-dimensional variable at different points in time).

In the static setting, with single dimensional types and single-dimensional allocations with payoff functions satisfying strictly increasing differences, a full characterization of incentive compatibility is immediate: A mechanism is incentive compatible if and only if the physical allocation is monotone and the transfers are pinned down by the revenue equivalence theorem.⁷

⁷If v is twice differentiable, single crossing differences is equivalent to $v_{x\theta}(x, \theta) > 0$ for all x, θ . Since $\int_r^{\theta} (v_\theta(x(s), s) - v_\theta(x(r), s)) ds = \int_{x(r)}^{x(\theta)} \int_r^{\theta} (v_{x\theta}(y, s) - v_\theta(y, s)) ds dy$, the claim follows.

Unfortunately it is not possible to find an equally attractive characterization for incentive compatibility in the dynamic model. In some sense, this is not too surprising. After all, it is well-known that characterizing incentive compatibility for models with multidimensional types (in this case, consider the types across different periods) is very difficult. At the same time, the fact that only the first period type is known to the agent at the moment of contracting lends additional structure to the model and it is possible to make progress even in the absence of tight necessary and sufficient conditions for the set of implementable allocations.

A very rough protocol for solving dynamic mechanism design problems can be given as follows. First, find the dynamic equivalent of the envelope formula (4) in the revenue equivalence theorem to compute the transfers as a function of the allocation process. Second, consider the relaxed principal's problem where her payoff is maximized subject to the constraint that the transfer is computed from (5). Third verify that the obtained solution satisfies the dynamic equivalent of the full incentive compatibility requirement (6).

3.2 Sequential Screening

We illustrate the general procedure now with a version of the sequential screening problem first investigated by Courty and Li (2000). An uninformed seller contracts with a privately informed buyer for two periods $t \in \{0, 1\}$ over the sale of a single indivisible good. At the beginning of the game, i.e. in $t = 0$, the buyer has some preliminary information about the value of the product. Let $\theta_0 \in \Theta_0 = [\underline{\theta}_0, \bar{\theta}_0]$ be her type representing this information. Let $F_0(\theta_0)$ denote the c.d.f. of the prior information on θ_0 and the corresponding density is denoted by $f_0(\theta_0)$. The buyer observes additional information on her true valuation as time proceeds. To capture this, we let the private information in period 1, $\theta_1 \in \Theta_1 = [\underline{\theta}_1, \bar{\theta}_1]$, denote the buyers' true willingness to pay for the object. This is without loss of generality since all information that can be considered in the allocation decision arrives in periods. Hence we just denote by θ_1 this total information available at the end of the game.

A key part of the model is the assumed dependence of θ_1 on θ_0 . If θ_0 is interpreted as the initial best estimate of the final θ_1 then it is reasonable to assume that θ_1 is increasing (in some sense) in θ_0 . We assume here that this dependence is in the sense of first order stochastic dominance (FOSD). To formalize this, we let $F(\theta_1 | \theta_0)$ denote the conditional distribution of θ_1 given θ_0 . Thus we assume:

$$\frac{\partial F(\theta_1 | \theta_0)}{\partial \theta_0} \leq 0 \text{ for all } \theta_0.$$

Observe that

$$\frac{\partial F(\underline{\theta}_1 | \theta_0)}{\partial \theta_0} = \frac{\partial F(\bar{\theta}_1 | \theta_0)}{\partial \theta_0} = 0,$$

since Θ_1 is the maximal support for θ_1 .

In order to apply the envelope theorem, we assume also that the conditional density $f(\theta_1 | \theta_0)$ is well defined, has full support, and that $|\partial f(\theta_1 | \theta_0) / \partial \theta_0|$ is uniformly bounded on $\Theta_0 \times \Theta_1$. Notice that $F_0(\theta_0)$ and $F(\theta_1 | \theta_0)$ induce jointly a distribution $F_1(\theta_1)$ on Θ_1 .

It is crucial to keep in mind that the moment of contracting in this model is $t = 0$, and hence at this stage θ_1 is not known to the buyer or the seller. The seller could always use static mechanisms based on either $F_0(\theta_0)$ or $F_1(\theta_1)$. The point of the sequential screening model is to show how the seller can do better than either of these alternatives. The idea is that the seller must leave information rent from the initial θ_0 to the buyer. The allocation can be made more efficient by using the information contained in θ_1 . If this efficiency gain can be realized without increasing the information rent to the buyer, then the seller gains relative to the two static alternatives. The model of sequential screening shows how this can be accomplished through a menu of *option contracts*.

Since we have defined θ_1 to be the true value of the object to the buyer, the simplest model of sales induces preferences over the probability of sales x and transfers p from the buyer to the seller as follows:

$$\begin{aligned} u_S(\theta_1, x, p) &= p, \\ u_B(\theta_1, x, p) &= \theta_1 x - p. \end{aligned}$$

Dynamic Direct Mechanisms and Incentive Compatibility

For this two period model we can then define a direct dynamic mechanisms as follows.

Definition 1 (Direct Dynamic Mechanism)

A direct dynamic mechanism is a pair of functions,

$$\begin{aligned} x &: \Theta_0 \times \Theta_1 \rightarrow [0, 1], \\ p &: \Theta_0 \times \Theta_1 \rightarrow \mathbb{R}_+. \end{aligned}$$

As in the example in the previous section, the buyer reports her type in each of the two periods. In the first, her report depends on her type θ_0 and in the second, it can depend on θ_0 , r_0 and on θ_1 . The dynamic revelation principle by Myerson (1986) can be invoked to show that there is no loss

of generality in concentrating on direct mechanisms where truth-telling is optimal on equilibrium path.⁸

We write

$$u(\theta_0, \theta_1) = \theta_1 x(\theta_0, \theta_1) - p(\theta_0, \theta_1)$$

for the final realized utility (or ex post utility), and

$$U(\theta_0; r_0) = \int_{\Theta_1} u(r_0, \theta_1) f(\theta_1 | \theta_0) d\theta_1$$

for the expected utility evaluated at the interim stage after period 0 reports. As in the static case, we let

$$V(\theta_0) = U(\theta_0; \theta_0).$$

Ex-post incentive compatibility then requires that for all reports r_1 in period $t = 1$,

$$u(\theta_0, \theta_1) \geq \theta_1 x(\theta_0, r_1) - p(\theta_0, r_1).$$

Interim incentive compatibility requires that for all r_0 ,

$$V(\theta_0) \geq U(\theta_0; r_0).$$

We say that the mechanism is *interim individually rational* if for all θ_0 ,

$$V(\theta_0) \geq 0.$$

3.2.1 Characterizing Incentive Compatible Mechanisms

We first present a useful characterization of incentive compatible mechanisms. The best case scenario would obviously be a result that gives an if and only if characterization for incentive compatible direct dynamic mechanisms. In the static case, with single dimensional types, such a characterization could be obtained under the additional assumption of payoffs that satisfy the Spence-Mirrlees single crossing condition. Recall that in this case, global incentive compatibility is equivalent to local incentive compatibility (or envelope formula) together with the monotonicity of the allocation rule in type.

In the current sales model, single crossing is clearly satisfied. We have also assumed that the types across the two periods are linked in the FOSD ordering. Hence there is a lot more structure

⁸In the current setting, it is also relatively easy to show that whenever truthful revelation of types is optimal on equilibrium path, it is also optimal off equilibrium path.

in place than for the *static two-dimensional* screening problem where little is known about the characterization of incentive compatibility. One should also bear in mind the simplifications that come from the sequential reporting structure.

We start with the easy part. For period $t = 1$, the first period report θ_0 is fixed and hence the incentive compatibility for reporting θ_1 reduces to the static incentive compatibility conditions.

Proposition 1 (Ex Post Incentive Compatibility)

A direct dynamic mechanism is incentive compatible with respect to the ex post type θ_1 if and only if:

1. (*Envelope Theorem*)

$$\frac{\partial u(\theta_0, \theta_1)}{\partial \theta_1} = x(\theta_0, \theta_1).$$

2. *Transfers are pinned down by the allocation and the payoff of the lowest type.*
3. (*Monotonicity*) $x(\theta_0, \theta_1)$ is non-decreasing in θ_1 for all θ_0, θ_1 .

One might hope that a similar proposition would hold for reporting the type θ_0 . Unfortunately this is not the case. Since the interim reports induce lotteries of payoffs (depending on the realization of θ_1), the static proof where incentive compatibility implies monotonicity fails (since now a similar monotonicity only needs to hold in expectation).⁹ While we will not get an if and only if statement for incentive compatibility, the envelope formula remains valid. Furthermore, monotonicity of the allocation rule is sufficient for full incentive compatibility. This means that the solution to the relaxed problem is the optimal mechanism if the allocation rule is monotone in both components.

Proposition 2 (Interim Revenue Equivalence)

In any incentive compatible direct dynamic mechanism:

1. (*Envelope Theorem*) For almost all θ_0 , we have:

$$V'(\theta_0) = - \int_{\Theta_1} x(\theta_0, \theta_1) \frac{\partial F(\theta_1 | \theta_0)}{\partial \theta_0} d\theta_1. \tag{7}$$

2. *Transfers are pinned down by the allocation and the payoff of the lowest type.*

⁹Krähmer and Strausz (2017) construct an equivalent static problem to each sequential screening problem of the above type.

3. Suppose that $x(\theta_0, \theta_1)$ is increasing in both components. Then there exists a transfer scheme $p(\theta_0, \theta_1)$ such that the direct dynamic mechanism (x, p) is incentive compatible.

Notice the factor $\partial F(\theta_1 | \theta_0) / \partial \theta_0$ in the first part of the proposition. It measures the impact of the first period type on the distribution of the second period type. Understanding how the initial type effects the information rents of the subsequent types is the key for all revenue maximizing dynamic mechanism design models. The second part in this proposition is just housekeeping. It is a result of the fact that the interim utilities can be evaluated by the envelope formula and directly through expected values from getting the object net of expected payment. Notice also that the first part implies immediately that any incentive compatible mechanism is individually rational if and only if $V(\underline{\theta}_0) \geq 0$.

3.2.2 Determining the Optimal Selling Mechanism

The next step is to determine the optimal mechanism. As in the static case, we express the payoff to the seller as the difference between the total surplus and the payoff to the buyer:

$$\int_{\Theta_0} \int_{\Theta_1} \theta_1 x(\theta_0, \theta_1) f(\theta_1 | \theta_0) f_0(\theta_0) d\theta_1 d\theta_0 - \int_{\Theta_0} V(\theta_0) f_0(\theta_0) d\theta_0.$$

Integrating by parts and using the envelope formulas allows us to write this as:

$$\int_{\Theta_0} \int_{\Theta_1} \left[\theta_1 + \frac{1 - F_0(\theta_0)}{f_0(\theta_0)} \frac{\partial F(\theta_1 | \theta_0)}{\partial \theta_0} \right] x(\theta_0, \theta_1) f(\theta_1 | \theta_0) f_0(\theta_0) d\theta_1 d\theta_0 - V(\underline{\theta}_0).$$

But now the solution of the relaxed problem where we ignore the global incentive compatibility constraints is close. Since we are using the Envelope formulas in computing the buyers' equilibrium payoff, we are making sure that the local incentive compatibility constraints hold. Clearly it is optimal to make the individual rationality constraint binding and set

$$V(\underline{\theta}_0) = 0.$$

Define a modified virtual value $\psi(\theta_0, \theta_1)$ by

$$\psi(\theta_0, \theta_1) \triangleq \theta_1 + \frac{1 - F_0(\theta_0)}{f_0(\theta_0)} \frac{\partial F(\theta_1 | \theta_0)}{\partial \theta_0}.$$

This modifies the classic Myersonian virtual value by multiplying the information rent component $(1 - F_0(\theta_0)) / f_0(\theta_0)$ by the factor

$$\frac{\partial F(\theta_1 | \theta_0)}{\partial \theta_0} f(\theta_1 | \theta_0),$$

measuring the impact of θ_0 on the distribution of θ_1 .

Since the value of the integral is linear in x , it is clearly optimal to set $x(\theta_0, \theta_1) = 1$ whenever $\psi(\theta_0, \theta_1) \geq 0$ in the relaxed program. If $\psi(\theta_0, \theta_1)$ is strictly increasing in both components, then this solution solves the revenue maximization problem. Hence we assume from now on that ψ is increasing in both arguments. To complete the description of the optimal mechanism, define the following function

$$q(\theta_0) = \min\{\theta_1 \in \Theta_1 \mid \psi(\theta_0, \theta_1) \geq 0\}.$$

Since ψ is increasing, $q(\cdot)$ is well defined. With the help of this function, we can characterize the optimal selling mechanisms.

Theorem 3 (Optimal Screening Mechanism)

If $\psi(\theta_0, \theta_1)$ is increasing in both arguments, then a direct dynamic mechanism (x, t) maximizes the seller's expected profit in the class of incentive compatible mechanisms if and only if

$$x(\theta_0, \theta_1) = \mathbb{I}_{\{\theta_1 \geq q(\theta_0)\}},$$

and the transfer is computed from the envelope formula.

Can we think of a nice indirect mechanism that implements this mechanism? We can consider option contracts for selling the good. In such a contract, the buyer buys for an up-front fee $p(\theta_0)$ the option of purchasing the good at strike price $q(\theta_0)$. Hence the mechanism seems to bear some relation to contracts that are actually observed in situations where uncertainty is gradually resolved and revealed about the value of the alternatives. Let $\{p(\theta_0), q(\theta_0)\}_{\theta_0 \in \Theta_0}$ denote the implementation of the optimal trading mechanism by these option contracts.

3.2.3 Information Rent and Orthogonalization

Esó and Szentes (2007) propose the following alternative way of approaching the problem. Rather than observing θ_1 , they let the buyer observe the *percentile*:

$$\gamma \triangleq F(\theta_1 \mid \theta_0).$$

Notice that γ is uniformly distributed on $[0, 1]$ for all θ_0 and thus γ is independent of θ_0 . This allows us to think of γ as the new information relative to θ_0 that is contained in θ_1 and the new variable γ is called the orthogonalized signal.¹⁰

¹⁰The use of orthogonalized signals also makes the analysis of the more complicated models of Pavan, Segal, and Toikka (2014) and Esó and Szentes (2017) more tractable.

The orthogonalized model allows for an easy comparison to the case where the seller also observes the orthogonalized signal. This is a useful benchmark that separates the initial private information from future realizations of symmetrically observed uncertainty. If the optimal allocation rule in the benchmark model satisfies monotonicity with respect to the orthogonal signal, the two models result in the same optimal mechanism. It should be noted though that the transfers implementing the optimal allocation can be quite different in the two models.

The orthogonalization procedure can also be used to see that the seller is always better off releasing the new information θ_1 rather than not releasing it. Li and Shi (2017) adds a word of caution regarding the use of this observation for general information management by the seller. They construct an example where the seller can achieve a higher expected revenue by showing a garbled version of θ_1 . Bergemann and Wambach (2015) extend the model of information design to consider sequential information release by the seller.

An important aspect in the argument of Esó and Szentes (2007) is that the choice of information disclosure is given directly in terms of the orthogonal information itself. Li and Shi (2017) show that if the information disclosure can be made contingent upon the realization of the initial private information, then the principal can frequently strictly improve his payoff by disclosing information only partially, and through a policy that is contingent upon the initial information θ_0 .

3.2.4 Multiple Buyers: Handicap Auctions

An early extension of the sequential screening model to cover multiple agents is the handicap auction in Esó and Szentes (2007). The model still has a single allocation stage and two reporting stages, but it allows for the presence of multiple bidders. Letting $\theta_0 \in \Theta_0$ and $\theta_1 \in \Theta_1$ denote the vectors of independent private value types $\theta_t^i \in \Theta_t^i$ of bidder i in t , we can then write a dynamic direct mechanism (x, p) as:

$$\begin{aligned} x &: \Theta_0 \times \Theta_1 \rightarrow \Delta^N, \\ p &: \Theta_0 \times \Theta_1 \rightarrow \mathbb{R}^N, \end{aligned}$$

where $\Delta^N = \{(x^1, \dots, x^N) \in \mathbb{R}_+^N \mid \sum_i x^i \leq 1\}$ is the set of possible allocations. The interpretation is that $x^i(\theta_0, \theta_1)$ gives the probability that i is assigned the object if the reported types are (θ_0, θ_1) .

Let

$$X^i(\theta_0^i, \theta_1^i) = \mathbb{E}_{\theta^{-i}} x^i(\theta^i, \theta^{-i})$$

and

$$P^i(\theta_0^i, \theta_1^i) = \mathbb{E}_{\theta^{-i}} p^i(\theta^i, \theta^{-i})$$

stand for the expected allocations and transfers of agent i . Thus, it is assumed that players other than player i report truthful their type when computing $X^i(\cdot)$ and $P^i(\cdot)$.

Similar to the static case, the analysis of incentive compatibility and revenue equivalence that we did in the case of a single buyer carries over to the case of many bidders with functions $(X^i(\theta_0^i, \theta_1^i), P^i(\theta_0^i, \theta_1^i))$. Hence we do not repeat the steps but go directly to the formula for expected revenue from bidder i in any incentive compatible mechanism:

$$\mathbb{E}_{\theta^i} P^i(\theta_0^i, \theta_1^i) = \int_{\Theta_0^i} \int_{\Theta_1^i} \psi^i(\theta_0^i, \theta_1^i) X^i(\theta_0^i, \theta_1^i) f^i(\theta_1^i | \theta_0^i) d\theta_1^i f^i(\theta_0^i) d\theta_0^i - U^i(\underline{\theta}_0^i),$$

where $\psi^i(\theta_0^i, \theta_1^i)$ is defined as before. Hence the total expected revenue is

$$\int_{\Theta_0} \int_{\Theta_1} \sum_{i=1}^N \psi^i(\theta_0^i, \theta_1^i) x^i(\theta_0, \theta_1) \prod_j f^j(\theta_1^j | \theta_0^j) d\theta_1 \prod_j f^j(\theta_0^j) d\theta_0 - \sum_{i=1}^N U^i(\underline{\theta}_0^i)$$

To maximize this expected revenue without violating the interim individual rationality constraints, it is clearly optimal to set $U^i(\underline{\theta}_0^i) = 0$ for all i , and to use the allocation rule:

$$x^i(\theta_0, \theta_1) = \begin{cases} 1, & \text{if } \psi^i(\theta_0^i, \theta_1^i) \geq \max\{0, \psi^j(\theta_0^j, \theta_1^j)\}, \text{ for all } j; \\ 0, & \text{if otherwise.} \end{cases}$$

The associated transfers can be computed from the envelope formula and binding participation constraint of the lowest type bidder.

3.3 Optimal Dynamic Mechanism

We now describe how the insights from the sequential screening environment can be extended to many periods and many allocations. We will be skipping a large number of details and for those we often refer the reader to Pavan, Segal, and Toikka (2014). We focus again mainly on the single agent case and consider a general discrete time horizon, $t \in \{0, 1, \dots, T\}$ with $T \leq \infty$.

We use the general dynamic model outlined in Section 2.1 with regularity assumptions that allow us to use the envelope theorem. The private information of the agent in period t is $\theta_t \in \Theta_t = [\underline{\theta}_t, \bar{\theta}_t]$. At the moment of contracting in $t = 0$, the agent knows θ_0 , and the stochastic process $F_t(\theta_{t+1} | \theta_t, x_t)$ governing $\{\theta_t\}_{t>0}$ but not the realizations θ_t for $t > 0$. The prior distribution on Θ_0 is given by

$F(\theta_0)$. We continue to assume that all distributions have full support and well defined conditional densities that are continuously differentiable with uniformly bounded derivatives (in t , x_t and θ_t).

A classic example of a stochastic process that satisfies these properties is the following $AR(1)$ process for $t \geq 1$:

$$\theta_t = \gamma\theta_{t-1} + \varepsilon_t, \tag{8}$$

where each ε_t is an i.i.d. draw from a single distribution $H(\cdot)$. This process has a moving average representation as:

$$\theta_t = \gamma^t\theta_0 + \sum_{s=0}^{t-1} \gamma^s \varepsilon_{t-s}. \tag{9}$$

This example is useful to keep in mind since it shows the general stochastic decay of the influence of the initial type θ_0 on subsequent types θ_t .

As in the previous section, we denote the agent's payoff from getting the allocation x_t in t is

$$\delta^t u_t(\theta_t, x_t, p_t) = \delta^t (v(x_t, \theta_t) - p_t).$$

We assume that $u_t(\cdot)$ is continuously differentiable and bounded in both arguments with uniformly bounded derivatives (in t, θ_t, x_t).

In a dynamic direct mechanism the buyer reports her type θ_t at each t . Letting $\theta^t = (\theta_0, \dots, \theta_t) \in \Theta^t$, we define a dynamic direct mechanism as $\{x_t, p_t\}_{t \geq 0}$:

$$x_t : \Theta^t \rightarrow X_t, \quad p_t : \Theta^t \rightarrow \mathbb{R}.$$

We say that the allocation rule is Markovian if $x_t(\theta^t) = x_t(\theta_t, x^{t-1})$. Any Markovian allocation rule x induces a Markov process whose transitions are given by

$$\theta_{t+1} \sim F_{t+1}(\cdot | \theta_t, x^t(\theta^t)),$$

and we denote this process by $\lambda[x]$. By the dynamic revelation principle in Myerson (1986), it is without loss of generality to consider a dynamic direct mechanism where the buyer reports her type $\theta_{i,t}$ truthfully in each t , and any such mechanism is said to be incentive compatible.

Canonical State Representation Similar to Section 3.2.3, it is possible to represent the process $\lambda[x]$ by a sequence of i.i.d. uniform random variables ε_t and a sequence of functions Z_t such that

$$Z_t(\theta_{t-1}, x^{t-1}, \varepsilon_t) \sim F_t(\cdot | \theta_{t-1}, x^{t-1}). \tag{10}$$

In this sense, we can view the process as being generated by θ_0 and the increments $(\varepsilon_t)_{t=1}^T$. The advantage of this new formulation is that now the stochastic increments ε_t are independent of θ_0 by construction. Pavan, Segal, and Toikka (2014) call this the *canonical representation* of the Markov process $(\theta_t)_{t=1}^T$. The reason for using this canonical representation of the Markov process, is that by the independence of ε_t relative to θ_0 , we can now vary θ_0 while keeping the future information structure fixed.

The *impulse response* at history (θ^t, x^{t-1}) relative to θ_0 is defined by:

$$I_t(\theta^t, x^{t-1}) \triangleq \frac{\partial Z_t(\theta_0, x^{t-1}, \varepsilon^t)}{\partial \theta_0}, \quad (11)$$

where ε^t satisfies for all $0 \leq s \leq t$:

$$\theta_s = Z_s(\theta_0, x^{s-1}, \varepsilon^s).$$

Similarly, we can define the intermediate impulse response function to previous shocks θ_s as:

$$I_{s,t}(\theta^t, x^{t-1}) \triangleq \frac{\partial Z_{s,t}(\theta_s, x^{t-1}, \varepsilon^t)}{\partial \theta_s}, \quad (12)$$

and by the chain rule we have:

$$\frac{\partial Z_t}{\partial \theta_0} = \prod_{s=1}^t \frac{\partial Z_{(s-1),s}}{\partial \theta_{s-1}}.$$

Using the canonical representation (10), we can evaluate the impulse response function (12) in terms of the original transition function:

$$I_t(\theta^t, x^{t-1}) = \prod_{\tau=0}^t \left(- \frac{\frac{\partial F_t(\theta_\tau | \theta_{\tau-1}, x^{\tau-1})}{\partial \theta_{\tau-1}^i}}{f_\tau(\theta_\tau | \theta_{\tau-1}, x^{\tau-1})} \right).$$

In order to make sure that the envelope theorem can be applied, we require that the canonical representation be differentiable and

$$|I_{t-1,t}(\theta^t, x^{t-1})| \leq B \text{ for some } B < \infty.^{11}$$

The impulse response function isolates the effects of the true type at t on the future evolution of $\tilde{\theta}_s$ for $s > t$. The other channel, i.e. through the dependence on x^t and therefore indirectly on θ^t depends only on the announced types in the mechanism and not the true types. The construction then proceeds by showing that the information rents are determined by the direct effects of payoffs. As in static allocation problem, the effects through announcements vanish by the first order condition for optimal announcements.

Payoff Equivalence We can then consider the equilibrium payoff to agent with type history θ^s :

$$V_s(\theta^s) = \mathbb{E} \left[\sum_{t=s}^T \delta^t (v_t(x_t(\theta^t), \theta_t) - p_t(\theta^t)) \right],$$

where the expectation is taken with respect to the continuation of the process $\lambda[x]$ from history $(\theta^s, x^s(\theta^s))$ onwards.

Theorem 4 (Local Incentive Compatibility)

If (x, p) is incentive compatible, then for all s, θ^s , $V_s(\theta^s)$ is Lipschitz continuous with derivative

$$V'_s(\theta^s) = \mathbb{E} \left[\sum_{t=s}^T I_{s,t}(\theta^t, x^{t-1}(\theta^{t-1})) \delta^t \frac{\partial v_t(x_t(\theta^t), \theta_t)}{\partial \theta_t} \right]. \quad (13)$$

The allocation rule thus pins down the incentive compatible payoff up to a constant as in the static case. To interpret the result, let

$$U(x, \theta) \triangleq \sum_{t=0}^T \delta^t v_t(x_t, \theta_t)$$

so that

$$\frac{\partial U}{\partial \theta_t} = \delta^t \frac{\partial v_t}{\partial \theta_t}.$$

The derivative of the indirect utility (13) then becomes:

$$V'_s(\theta^s) = \mathbb{E} \sum_{t=s}^T I_{s,t} \frac{\partial U}{\partial \theta_t}.$$

The impulse response function measures the effect of a small change in θ_s on θ_t and $\frac{\partial U}{\partial \theta_t}$ measures the induced change in period t utility. The transfers that support the indirect utility of the agent can then be derived just as in static model.

As discussed earlier, sufficient conditions for the optimality of the candidate mechanism given by the local incentive conditions are much more difficult to obtain in the dynamic setting than in the static setting. Even though θ_t is one-dimensional at each t , the resulting allocation is multi-dimensional since reports in period t affect the allocation in all future periods. Pavan, Segal, and Toikka (2014) offer a necessary and sufficient condition for the implementability of the allocation process that corresponds almost exactly to the full incentive compatibility condition 6 of Section 3.1. While the condition itself is not easy to verify, it suggests stronger notions that are easier to verify. Their condition of strong monotonicity that holds whenever the type evolution is independent of past allocations and all future allocations are increasing in all reported types turns out to be particularly useful in applications.

Relaxed solution Consider next the optimization problem of a principal that has her own dynamic payoffs given by

$$\sum_{t=0}^T \delta^t (p_t - c_t(x_t)).$$

She designs a mechanism to maximize her own payoff. As always, we can write the principal's payoff as the difference between the social surplus and the agent's information rent. Hence the problem is equivalent to

$$\begin{aligned} \max_{(x,p)} \mathbb{E}^x \sum_{t=0}^T \delta^t (v_t(x_t(\theta^t), \theta_t) - c_t(x_t(\theta^t))) \\ - \mathbb{E}^x \frac{1 - G(\theta_0)}{g(\theta_0)} \left[\sum_{t=0}^T \delta^t I_t(\theta^t, x^{t-1}(\theta^{t-1})) \frac{\partial v_t(x_t(\theta^t), \theta_t)}{\partial \theta_t} \right] - V_0(\underline{\theta}_0). \end{aligned} \quad (14)$$

subject to the incentive compatibility conditions and the period 0 participation constraints:

$$V_0(\theta_0) \geq 0.$$

We denote the first line in the objective function by $\mathbb{E}^x [S(x, \theta)]$ to represent the social surplus. We have built the local incentive compatibility conditions into the objective function by using the envelope formula to represent the buyer's information rent. Typically, the individual participation constraint will bind at the optimum for the lowest type and thus $V_0(\underline{\theta}_0) = 0$. Solving

$$\max_x \mathbb{E}^x \left[S(x, \theta) - \frac{1 - G(\theta_0)}{g(\theta_0)} \sum_{t=0}^T \delta^t I_t \frac{\partial v_t}{\partial \theta_t} \right]$$

involves dynamic programming and is not easy in general. If the process of $(\theta_t)_{t=1}^T$ does not depend on the allocation x^t and there are no intertemporal restrictions on x_t , then a pointwise solution is often possible. If the relaxed problem allows for an explicit solution, one can check if the sufficient conditions for full incentive compatibility are satisfied. For the examples that we describe below, the solution of the relaxed problem can be characterized in sufficient detail to allow us to verify sufficient conditions for full incentive compatibility.

If the solution of the relaxed problem is not fully incentive compatible, little is known about the methods for solving the problem. Even if a full solution is not possible, the necessary local incentive compatibility conditions as expressed in the envelope formula may give some qualitative insights about the properties of the optimal dynamic contract.

3.4 Implications and Applications of Optimal Mechanisms

We now discuss two notable application of the optimal mechanism to specific auction and contracting environments: (i) intertemporal licensing and (ii) irreversible sale of durable good.

3.4.1 Bandit Auctions

A single indivisible object is allocated in each period amongst n possible bidders that learn about their true valuation for the good. The type of bidder i changes only in periods t where she is allocated the good: if $x_t^i = 0$, then $\theta_{t+1}^i = \theta_t^i$, if $x_t^i = 1$, then

$$\theta_{t+1}^i = \theta_t^i + \varepsilon^i (n^i(t)) \tag{15}$$

where ε^i is a random variable whose distribution depends on the number of periods up to t , $n^i(t)$, in which the good has been allocated to i . For some stochastic processes such as the normal learning process outlined in Section 2.1.2, the number of observations from the process (here $n^i(t)$) and the current posterior mean (here θ_t^i) form a sufficient statistic. We can interpret the allocation process as intertemporal licensing where the current use of the object is determined on the past and current reports of the bidders. Notably, the assignment of the object can move back and forth between the bidders as a function of their reports. Pavan, Segal, and Toikka (2014) and Bergemann and Strack (2015) consider a revenue maximizing auction for the special case of the multi-armed bandit model in discrete or continuous time, respectively. The efficient allocation policy under private information was analyzed earlier in Bergemann and Välimäki (2010).

A useful aspect of the bandit model with the additive noise model is the easily verified property that:

$$\prod_{\tau=t}^s \left(- \frac{\frac{\partial F_t^i(\theta_{\tau+1}^i | \theta_\tau^i)}{\partial \theta_\tau^i}}{f_\tau^i(\theta_{\tau+1}^i | \theta_\tau^i)} \right) = 1. \tag{16}$$

Hence the revenue maximization problem is now turned (again using the usual steps) into a modified bandit problem where the seller maximizes

$$\max_{x \in X} \mathbb{E} \sum_{t=0}^T \sum_{i=1}^N \delta^t \left[\theta_t^i - \frac{1 - F^i(\theta_0^i)}{F^i(\theta_0^i)} \right] x_t^i(\theta),$$

where $X = \{(x^1, \dots, x^N) \in \mathbb{R}_+^N | \sum_i x^i = 1\}$. Stated in this form, the problem can be solved using the dynamic allocation index, the Gittins index. Pavan, Segal, and Toikka (2014) verify that

the solution satisfies the average monotonicity condition and is hence implementable. Thus, the resulting dynamic optimal auction proceeds by finding the bidder with the highest valuation after taking into account the handicap that is determined exclusively by the initial private information θ_0^i . Moreover, by (16), the impulse response function, and hence the handicap is constant in time and determined only through the initial shock.

Kakade, Lobel, and Nazerzadeh (2013) consider a class of dynamic allocation problems that includes the above Bandit problem. By imposing a separability condition (additive or multiplicative) on the interaction of the initial private information and all subsequent signals, they obtain an explicit characterization of the revenue-maximizing contract and derive transparent sufficient conditions for the optimal contract.

Bergemann and Strack (2015) consider general time-separable allocation problems in continuous time. By restricting attention to problems where (i) the set of feasible allocations at time t is independent of the history of the allocations, and (ii) the flow utility functions depends only on the initial and current private information, they can leverage the structure of the problem to frequently obtain closed-form solutions of the optimal contract. In the leading example of repeat sales of a good or service, they establish that the commonly observed contract features such as flat rates, free consumption units and two-part tariffs emerge as part of the optimal contract.

3.4.2 Selling Options

The second illustrative example is a stopping problem rather than a recurrent allocation problem. Suppose we would like to allocate a single object among N bidders, but we can allocate it only once and for all. Thus, the seller faces a stopping problem, and at the moment of stopping, he needs to make a decision to whom to allocate the object. Suppose the evolution of the willingness to pay by bidder i is given by:

$$\theta_t^i = \gamma\theta_{t-1}^i + \varepsilon_t^i,$$

with $\theta_0^i \sim G^i(\theta_0^i)$, $\varepsilon_t^i \sim H^i(\cdot)$, i.i.d. If we set $\gamma = 1$, we are essentially dealing with the model of Board (2007).

We can now compute the indirect utility function in the familiar way,

$$V_0^i(\theta_0^i) = \mathbb{E} \sum_{t=0}^T \delta^t \frac{1 - G^i(\theta_0^i)}{g^i(\theta_0^i)} \gamma^t x_t^i(\theta),$$

and find that the expected revenue to the seller is

$$\mathbb{E} \sum_{t=0}^T \sum_{i=1}^N \delta^t \left[\theta_t^i - \frac{1 - G^i(\theta_0^i)}{g^i(\theta_0^i)} \gamma^t \right] x_t^i(\theta),$$

The seller's problem is thus an optimal stopping problem, and her decision in period t is whether to stop the process and collect

$$\max_i \left\{ \theta_t^i - \frac{1 - G^i(\theta_0^i)}{g^i(\theta_0^i)} \gamma^t \right\},$$

or to continue until $t + 1$ and draw a new valuation vector $\theta_{t+1} = \gamma\theta_t + \varepsilon_t$ for the bidders. As time progresses and t increases, the distortion relative to the planner's solution in the allocation diminishes.

3.5 The Role of the Markovian Environment

A feature common to almost all of the above contributions is that the private information of the agent is represented by the current state of a one-dimensional Markov process, and that the new information that the agent receives is controlled by the current state, and in turn, leads to a new state of the Markov process. Now, in any model where the initial state is the current state of a recurrent Markov process, such as in Battaglini (2005), the informativeness of the initial state about future states is vanishing over time. With the impulse response of the initial state vanishing over time this then implies that the allocative distortion vanishes in the long-run.

By contrast, a number of recent contributions considered the possibility that the initial private information is about a parameter of the stochastic process itself, such as the drift or the variance of the process. For example, Boleslavsky and Said (2013) let the initial private information of the agent be the mean of a multiplicative random walk. This changes the impact that the initial private information has on the future allocations. The distortions in the future allocation may now increase over time rather than decline as in much of the earlier literature. The reason is that the influence of the parameter of the stochastic process on the valuation may increase over time. Pavan, Segal, and Toikka (2014) and Skrzypacz and Toikka (2015) report similar findings.¹²

The impulse response function in Boleslavsky and Said (2013) involves the number of past realized upticks and downticks of the binary random walk. Bergemann and Strack (2015) consider

¹²This is equivalent to assuming that the private information of the agent corresponds to the state of a two-dimensional Markov process, whose first component is constant after time zero, but influences the transitions of the second component.

the continuous time version of the multiplicative random walk, the geometric Brownian motion. Interestingly, in the continuous time version, the impulse response function is simply the expected number of upticks or downticks, which is a deterministic function of time and the initial state. Correspondingly, the handicap factor is increasing linearly over time, and the optimal contract prescribes a deterministic time at which the trade ends, thus suggesting a leasing contract with fixed term length. More generally, Bergemann and Strack (2015) allow the valuation process of the buyer to be either the arithmetic, geometric, or mean-reverting Brownian motion. Across these classes of models, they show the allocative distortion of the revenue-maximizing contract can be constant, decreasing, increasing or even random over time depending on the precise nature of the private information.

4 Dynamic Populations

In this section, we consider mechanism design problems where the population of privately informed agents changes over time. To fix ideas, we still assume that the mechanism designer is a seller who has a fixed capacity of objects for sale. Her problem is now to find an incentive compatible mechanism that maximizes her expected revenue. The model presented below can be viewed as a generalization of the revenue management model in operations research to cover incomplete information.

It is possible to distinguish two types of problems within this setting. In the first, the process of arrivals for potential buyers is known at the outset. With independent private types, this model is quite close to the applications surveyed in the last section. The dynamic pivot mechanism then offers an individually rational and incentive compatible implementation of the efficient allocation rule. For the dynamic revenue maximization problems, dynamic stopping problems based on virtual valuations of the buyers provide optimal solutions.

In the second class of models, the distribution of valuations or the process of arrivals of potential buyers is itself unknown at the outset. The mechanism designer (and the agents in the mechanism) learn more about the distribution of future demand from past realizations of arrivals and valuations for the buyers. An early example of a model of this type with unknown value distributions in the static case is Segal (2003). Gershkov and Moldovanu (2009b) points out problems that arise for models of this type when agents arrive over time and when payments have to be made at the time of allocation decision. We illustrate the problems that arise with a slight modification of the initial

example that allows for correlated types.

Example: Continued with Correlation We continue to consider an auction of a single indivisible object between two bidders $i \in \{0, 1\}$. But now bidder i is present in the mechanism only in period $t = i$. Hence the possible social allocation decisions are given by $x \in \{0, 1\}$, where the first choice indicates allocating the object to bidder 0 in period 0 and the second indicates allocating to bidder 1 in period 1.

The valuation of bidder 0 is $\theta_0 \in [0, 1]$. Bidder 1 learns her value in period 1, and $\theta_1 \in \{0, 1\}$. The prior probability on the high type is p and we further assume that the types of the two agents are correlated as follows:

$$\Pr\{\theta_1 = 1 \mid \theta_0\} = \frac{1}{3}, \quad \text{if } \theta_0 < \frac{1}{2},$$

and

$$\Pr\{\theta_1 = 1 \mid \theta_0\} = \frac{2}{3}, \quad \text{if } \theta_0 \geq \frac{1}{2}.$$

Since agent 0 is present in $t = 0$ only, the optimal allocation is determined solely by type θ_0 of bidder 0 and it is given by:

$$x^*(\theta_0) = 0 \Leftrightarrow \theta_0 \in \left[\frac{1}{3}, \frac{1}{2}\right) \cup \left[\frac{2}{3}, 1\right].^{13}$$

Notice that this allocation rule is not monotone in θ_0 and hence not incentive compatible for any possible transfer rule $p(\theta_0)$. Hence we conclude that incentive compatibility and efficiency are not mutually possible. If we allowed agent 0 to receive and make transfers in $t = 1$, efficiency would be restored by reverting to the usual second price auction allocation and transfer rules:

$$p_0(\theta_0, \theta_1) = (1 - x^*(\theta_0))\theta_1, \quad p_1(\theta_0, \theta_1) = x^*(\theta_0)\theta_0.$$

Importantly, this example also illustrates why achieving dominant strategies is not possible to achieve in dynamic mechanisms even with unrestricted transfers. If agent 1 always reported $\theta_1 = 1$, then the above mechanism would not be incentive compatible for agent 0. Since agent 1 has no private information in $t = 0$, the requirement of periodically ex post incentive compatibility is trivially satisfied.

Finally notice that the notion of marginal contribution of agent 0 is ambiguous in this case. When computing the contribution of agent 0 to the social welfare, it is not clear if we should use the expected value if allocating to 1 with the private information of agent 0 or without that information (i.e. if only the allocational possibilities or also the information should be removed when considering the social welfare when 0 is not present).

Short vs. Long Lived Bidders; Observable vs. Unobservable Arrival Further distinctions between the models are possible along several dimensions. The agents may be long-lived or short lived as in many revenue management models. Gershkov and Moldovanu (2009a) extend the classic revenue management models a la Gallego and van Ryzin (1994) to cover more general allocation models with short-lived buyers. If they are long-lived, models may distinguish between that case where their arrival is observable to the mechanism designer and where it is not. In the latter case, the agents also maximize over the time at which they report their arrival (and type). Board and Skrzypacz (2015) cover the case of stochastically arriving buyers that can time their purchases. Gershkov, Moldovanu, and Strack (2017) develops a revenue maximizing model where the arrival rates of the buyers over time are initially uncertain and must be learned over time.

The key to solving this type of revenue maximization problems is to find the appropriate externality payment mechanism for the problem at hand where virtual valuations are used instead of the true valuations. Hence the exercises are closely related in the spirit to the classical Myerson auctions and the techniques are familiar from the previous section of this survey. The key first step is to express the payoff to the agents using an envelope formula and then make sure that the model has enough monotonicity to ensure full incentive compatibility of the solution of the relaxed problem. In the case where the agents time their reporting to the mechanism designer strategically, an additional constraint of monotonicity with respect to reporting time must also be verified. It is interesting to note that strategic timing by the agents ends up enhancing the revenue accruing to the seller in the optimal mechanisms in Board and Skrzypacz (2015) and Gershkov, Moldovanu, and Strack (2017).

Finally, Garrett (2016) considers a model where the type of the agent changes stochastically, providing an additional reason for the strategic timing of reporting one's arrival. By analyzing the optimal commitment paths for prices, Garrett (2016) provides a new rationale for fluctuating prices even in otherwise completely stationary models.

5 Connections to Nearby Models

In this section, we discuss briefly two classes of dynamic contracting models that do not assume quasi-linear payoffs. Since Rogerson (1985), models of dynamic moral hazard have discussed the smoothing of dynamic risks in models with incentive problems. In dynamic settings, the distinction between dynamic moral hazard and adverse selection is almost impossible to make and many models

that share the informational structure with our general dynamic model have been discussed under the name of dynamic moral hazard (see for example the early formulations in Thomas and Worrall (1988), and Phelan and Townsend (1991)). The key difference between these models and those discussed in the previous sections is that with risk averse preferences, the trade-off between efficient physical allocation and efficient risk allocation emerges. In the first subsection, we discuss briefly the models employed in new public finance from this viewpoint.

In models of financial economics, a key assumption is that the privately informed managers do not have sufficient funds to buy the entire enterprise. This is typically formalized through a limited liability constraint stating that managers (the agent) cannot make payoffs to owners (the principal). Recent work starting with Clementi and Hopenhayn (2006), DeMarzo and Sannikov (2006), DeMarzo and Fishman (2007), and Biais, Mariotti, Plantin, and Rochet (2007) has analyzed the problem of incentivizing a manager that observes privately the cash flow of a firm. In the second subsection, we describe how the model and the results in this model compare to the results from the general dynamic model in this survey.

5.1 New Public Finance: Risk-Averse Agent

In most mechanism design problems, the key problem for the designer can be formulated as follows: what is the most advantageous way of providing the agent with a fixed level of utility v . With risk-averse agents and a risk-neutral principal, optimal contracts provide some amount of insurance, but incentive compatibility precludes the possibility of full insurance.

In dynamic problems, the principle of dynamic programming suggests that this question should be decomposed into current incentives and future incentives: conditional on each report θ_t , the current allocation $(x(\theta^t), p(\theta^t))$ is decided and a continuation payoff $w(\theta^t)$ is induced for the future. The idea is to give relatively high consumption to the agent in the periods where marginal utility from consumption is high, but to preserve incentive compatibility, this must come at the cost of lower continuation payoffs. Spear and Srivastava (1987) were the first to use this recursive formulation for dynamic incentive provision: all the agent needs to know about the consequences of her reports at each point in time are their payoff consequences and by the principle of dynamic programming, these are summarized in $(x(\theta^t), p(\theta^t), w(\theta^t))$. Of course for the dynamic contract to be incentive compatible and consistent, two conditions are needed:

$$u(x(\theta_t, \theta^{t-1}), p(\theta_t, \theta^{t-1}), \theta_t) + \delta w(\theta_t, \theta^{t-1}) \geq u(x(\theta'_t, \theta^{t-1}), p(\theta'_t, \theta^{t-1}), \theta_t) + \delta w(\theta'_t, \theta^{t-1})$$

for all t, θ_t, θ'_t ; and

$$w(\theta^{t-1}) = \mathbb{E}_{\theta^t} [u(x(\theta_t, \theta^{t-1}), p(\theta_t, \theta^{t-1}), \theta_t) + \delta w(\theta_t, \theta^{t-1}) | \theta^{t-1}, x^{t-1}],$$

for all t, θ^{t-1} .

The first of these constraints is the usual incentive compatibility constraint that we encountered earlier in sequence form with the quasi-linear utilities and the second one is called the promise keeping constraint that makes sure that the promised utilities are realized through the allocation decisions for truthful reports.

The literature on optimal risk allocation over time has focused on two types of results. The first is a quite striking and initially counterintuitive finding. Consider a risk-neutral principal with a linear cost of providing consumption $x_t \geq 0$ to a privately informed agent. The objective of the principal is to minimize the cost of providing an initial payoff of v to the agent. The periodic income of the agent is her private information θ_t , her type, which is assumed to be an i.i.d. draw $\theta_t \sim F(\cdot)$ in all periods. Her payoff is given by a continuous and strictly concave utility function $u(x)$ that is unbounded from below. The famous immiseration result by Thomas and Worrall (1990) states that along the optimal consumption path, we have almost surely $x_t \rightarrow 0$ and therefore $u(x_t) \rightarrow -\infty$. This result is also shown by Atkinson and Lucas (1992) to hold for a utilitarian social planner operating subject to an aggregate feasibility constraint.¹⁴

To understand this result, note that the concavity in the utility function makes the threat of low future consumptions an effective way of delivering incentives for current reports. A low current θ_t is associated with a relatively high current consumption x_t , and to preserve incentive compatibility, this must be associated with a relatively low future payoff. A simple variational argument establishes that the marginal cost for the principal for providing promised utility v must be a martingale. The martingale convergence theorem establishes the convergence of this marginal cost. Promised utilities cannot converge to a finite constant since this would result in optimal full insurance within period contradicting incentive compatibility (with constant future promises) and the result follows.

The second key result is the inverse Euler equation for the provision of consumption. Rogerson (1985) considers the optimal contracting problem between a risk-neutral principal and a risk-averse agent choosing an unobservable action (and no private information). One of the key findings in

¹⁴Feasible allocations are constrained by the endowment process in an exchange economy.

that paper is that the optimal form of providing dynamic incentives to the agent takes the form:

$$\frac{1}{u'(x_t)} = \frac{1}{\delta \mathbb{E}u'(x_{t+1})}.$$

Notice that the terms in this formula are the inverses of the lhs and the rhs of the usual Euler equation and hence the name. This same characterization holds for dynamic models of incomplete information where $x(\theta^t)$ is determined based on the reported types as long as current utility does not depend on θ_t . In the literature on dynamic public finance, monetary transfers are not present, and the allocation takes typically the form $x(\theta^t) = (c(\theta^t), l(\theta^t))$, where c denotes the consumption and l denotes the labor supply. If $u(x, \theta)$ takes the form

$$u(x, \theta) = v(x) - g(l, \theta),$$

the inverse Euler equation for the consumption allocations holds as shown in Golosov, Kocherlakota, and Tsyvinski (2003). By Jensen's inequality, one sees immediately that if the agent is allowed to save at interest rate $1/\delta$, she will optimally do at the consumption path $x(\theta^t)$. This observation has given rise to the large literature on optimal contracting with hidden savings. (We should include some references here).

The immiseration results have led to a reconsideration of the assumption of full commitment on the two sides of the contracting problem. Assuming that the outside option of the agent is increasing in her current type leads to a new qualitative feature in the optimal contract. As observed by Harris and Holmstrom (1982) and Kocherlakota (1996), without commitment, the promised utility in the contract moves upward over time as higher and higher types are reached.

An alternative approach is taken in Farhi and Werning (2007), where the planner discounts future periods less than the generation born in the initial period. This consideration for intergenerational distribution leads to solutions that often generate optimal income distributions that do not converge towards the immiseration outcome.

If one assumes that the set of possible types is a connected interval and that the process of types has full support, then the Bellman equation of the agent can be written as:

$$V'(\theta_t) = \frac{\partial u(x_t(\theta^t), \theta_t)}{\partial \theta_t} + \int_{\underline{\theta}}^{\bar{\theta}} V'(\theta_{t+1}) \frac{\partial f(\theta_{t+1} | \theta_t)}{\partial \theta_t} d\theta_{t+1}.$$

Integration by parts gives:

$$\begin{aligned} V'(\theta_t) &= \frac{\partial u(x_t(\theta^t), \theta_t)}{\partial \theta_t} + \mathbb{E} \left[\frac{-\frac{\partial F(\theta_{t+1} | \theta_t)}{\partial \theta_t}}{f(\theta_{t+1} | \theta_t)} V'(\theta_{t+1}) \right] \\ &= \frac{\partial u(x_t(\theta^t), \theta_t)}{\partial \theta_t} + \mathbb{E} [I_{t,t+1}(\theta^t) V'(\theta_{t+1})]. \end{aligned}$$

In other words, a similar envelope theorem characterization for the agent's utility is still possible in this model. The next step of substituting the agent's payoff into the principal's objective unfortunately fails because of the lack of quasi-linear utility. As a result, solving the model is in general more difficult than in the quasi-linear case and numerical methods are typically needed. This also implies that checking full incentive compatibility becomes much harder in this class of models.

5.2 Limited Liability

In financial economics, a key incentive problem is between a privately informed manager (agent) and an uninformed owner (principal). Limited liability protection on part of the agent implies an upper bound on the transfers that can be made from the agent to the principal. Often this constraint takes the form that all transfers must be from the principal to the agent. This prevents of course the principal from selling the enterprise to the agent at the outset. Since recent surveys of this large literature exist (see for example Biais, Mariotti, and Rochet (2013)), we do not attempt a comprehensive survey of the topic.

A recent paper Krasikov and Lamba (2016) takes up a particular quasilinear version of our general dynamic model with the added feature of limited liability. In a two-state Markov chain model of information for their model, they consider optimal contracting between a principal and a manager that knows her private cost type. The restriction to a binary type set allows Krasikov and Lamba (2016) to solve fully the optimal contract using techniques similar to Battaglini (2005). Echoing the results from earlier papers in dynamic mechanism design, the paper shows that the optimal contract converges eventually to the efficient static contract. Prior to the path of reaching this region, the contract displays various types of inefficiencies reflecting the limited liability constraints as well as distortions arising from the principal's attempt to extract information rent from the agent.

We believe that a lot of work still remains in the analysis of the risk-neutral setting with additional contractual limitations imposed on the problems. Partial verifiability as in the case of a manager that cannot exaggerate the cash flow of her firm and limited liability constraints are good starting points for this analysis. More work is also needed towards a better understanding of ex post participation constraints.

5.3 Bounding the Performance of Mechanisms

The intertemporal allocations and commitments that resulted from the dynamic mechanism balanced trade-offs over time. These trade-offs were based on the expectations of the agents and the principal over the future states. In this sense, all of the mechanisms considered were Bayesian solutions and relied on a shared and common prior of all participating players. Yet, this clearly is a strong assumption and a natural question would be to what extent weaker informational assumptions, and corresponding solution concepts, could provide new insights to the format of dynamic mechanisms. For example, the sponsored search auctions which provide much of the revenue for the search engines on the web, are clearly repeated and dynamic allocations with private information, yet, most of the allocations and transfer are determined by spot markets rather than long-term contracts. An important question then is why not more transactions are governed by long-term arrangements that could presumably share the efficiency gains from less distortionary allocations between the buyers and the seller. An important friction to long-term arrangements is presumably the diversity in expectations about future events between buyer and seller. In a recent paper, Mirrokni, Leme, Tang, and Zuo (2017) provide lower bounds for a revenue maximizing mechanism in which the players do not have to agree on their future expectations. The mechanism that achieves the lower bound in fact satisfies the interim participation and incentive constraints for all possible realizations of future states. This approach reflects the recent interest of theoretical computer science in dynamic mechanism design, see for example Papadimitriou, Pierrakos, Psomas, and Rubinstein (2016), and mechanism design more generally. But in contrast to the approach most commonly taken by economic theory who attempt to explicitly identify and design the optimal mechanism, theoretical computer science often describes achievable bounds on the performance. The bounds are frequently achieved by mechanisms that have computational advantages in terms of computational complexity relative to the, possibly unknown, exact optimal mechanism.

6 Concluding Remarks

It was our objective to present some of the recent work on dynamic mechanism design. We hope we have conveyed the scope and the progress that has been made in the past decade. Our discussion should have also indicated that many interesting questions remain wide open. We shall describe some of them in these final remarks.

We argued that the optimal dynamic mechanism allows the seller to receive a large share of the

surplus by replacing the sequence of periodic, and hence static and ex post participation constraints, by a single ex ante participation constraints. We discussed above that an important friction to long-term arrangements is presumably the diversity in expectations about future events among the players. Here, a natural direction is to ask to what extent the insights from static mechanism design can be transferred to dynamic settings. Mookherjee and Reichelstein (1992) establish that in static environments the revenue maximizing allocation can frequently be implemented by dominant rather than Bayesian incentive compatible strategies. Similarly, Bergemann and Morris (2005) present conditions for static social choice functions under which an allocation can be implemented for all possible interim beliefs that the agents may hold. The robustness to private information is arguably an even more important consideration in dynamic environments.

The central problem that the literature of dynamic mechanism has addressed is how to provide incentives to report the sequentially arriving private information. Thus, the central constraints on the design are given by the sequence of interim incentive compatibility conditions. The participation constraints on the other hand have, perhaps surprisingly for a dynamic environment, have received much less attention. Now, sometimes a dynamic mechanism can guarantee the ex-ante participation constraints as well as the interim (or periodic) ex-post constraints. The dynamic pivot mechanism that governed the dynamically efficient allocation provided such an instance. By contrast, the dynamic revenue maximization contract only imposed the participation constraint at time zero, thus ex-ante relative to all future arriving private information. In particular, the dynamic mechanism cannot provide any guarantees about ex post participation constraints. Indeed, Krämer and Strausz (2015) show that sequential screening frequently reduces to a static screening solution if the seller has to meet the ex post rather than the ex ante participation constraints of the buyers. Moreover, if the dynamic mechanism improves upon a static mechanism in the sequential screening model, then the ex post participation constraint severely limits the ability of the seller to extract surplus through option contracts as shown in Bergemann, Castro, and Weintraub (2017).

A related but distinct issue that is likely to require more analysis is the timing of the contractual agreement between principal and agents. Much of the current analysis assumes that the arrival of the agents is known to the principal and that the principal can make a single, take-it-or-leave-it offer at the moment of the agent's arrival. This constraint, while natural in a static setting, is much less plausible in dynamic settings. In particular, it explicitly excludes the possibility for the agent to postpone and delay the acceptance decision to a later time when he may have additional information about the value of the contract offered to him. A notable exception is Garrett (2016) who allows

for the random arrival of buyers and suggests that it may offer an explanation for temporary price reduction even in an otherwise stationary environment. The participation constraints may impose important constraints on the principal either because the agents may arrive over time and the arrival time may be private information to the agent, or because the principal lacks in commitment power relative to future offers. For example, Lobel and Paes Leme (2017) assume that the seller can make future commitments to the current buyer, yet the seller is unable to commit to future offers that he will make to newly arriving buyers. The resulting inability to preclude future, and possibly more favorable offers, presents a new option to the current buyer, and hence strengthens his bargaining position, and limits the ability of the seller to extract surplus. Bergemann and Strack (2017) analyze the dynamic revenue maximizing contract subject to the restriction that the seller has to offer a stationary contract. Thus the seller is forced to renew the contract offer in every period, both to newly arriving and currently waiting customers. They show that the ability to postpone the acceptance of an offer to a future period can increase the value of the buyer and can lead to a more efficient allocation resulting in equilibrium.

References

- ATHEY, S., AND I. SEGAL (2013): “An Efficient Dynamic Mechanism,” *Econometrica*, 81, 2463–2485.
- ATKENSON, A., AND R. LUCAS (1992): “On the Efficient Distribution with Private Information,” *Review of Economic Studies*, 59, 427–454.
- BARON, D., AND D. BESANKO (1984): “Regulation and Information in a Continuing Relationship,” *Information Economics and Policy*, 1, 267–302.
- BATTAGLINI, M. (2005): “Long-Term Contracting with Markovian Consumers,” *American Economic Review*, 95, 637–658.
- BERGEMANN, D., F. CASTRO, AND G. WEINTRAUB (2017): “The Scope of Sequential Screening with Ex-Post Participation Constraints,” in *18th ACM Conference on Economics and Computation*.
- BERGEMANN, D., AND S. MORRIS (2005): “Robust Mechanism Design,” *Econometrica*, 73, 1771–1813.
- BERGEMANN, D., AND A. PAVAN (2015): “Introduction to Symposium on Dynamic Contracts and Mechanism Design,” *Journal of Economic Theory*, 159, 679–711.
- BERGEMANN, D., AND M. SAID (2011): “Dynamic Auctions: A Survey,” in *Wiley Encyclopedia of Operations Research and Management Science*, ed. by J. Cochran, no. 1757, pp. 1511–1522. Wiley, New York.
- BERGEMANN, D., AND P. STRACK (2015): “Dynamic Revenue Maximization: A Continuous Time Approach,” *Journal of Economic Theory*, 159, 819–853.
- (2017): “Stationary Dynamic Contracts,” Discussion paper, Yale University and UC Berkeley.
- BERGEMANN, D., AND J. VÄLIMÄKI (2002): “Information Acquisition and Efficient Mechanism Design,” *Econometrica*, 70, 1007–1033.
- (2010): “The Dynamic Pivot Mechanism,” *Econometrica*, 78, 771–790.

- BERGEMANN, D., AND A. WAMBACH (2015): “Sequential Information Disclosure in Auctions,” *Journal of Economic Theory*, 159, 1074–1095.
- BESANKO, D. (1985): “Multi-Period Contracts Between Principal and Agent with Adverse Selection,” *Economics Letters*, 17, 33–37.
- BIAIS, B., T. MARIOTTI, G. PLANTIN, AND J. ROCHET (2007): “Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications,” *Review of Economic Studies*, 74, 345–390.
- BIAIS, B., T. MARIOTTI, AND J.-C. ROCHET (2013): “Dynamic Financial Contracting,” in *Advances in Economics and Econometrics, Tenth World Congress*, ed. by D. Acemoglu, M. Arellano, and E. Dekel, vol. 1. Cambridge University Press.
- BOARD, S. (2007): “Selling Options,” *Journal of Economic Theory*, 136, 324–340.
- BOARD, S., AND A. SKRZYPACZ (2015): “Revenue Management with Forward-Looking Buyers,” *Journal of Political Economy*, 124, 1046–1087.
- BOLESZLAVSKY, R., AND M. SAID (2013): “Progressive Screening: Long-Term Contracting with a Privately Known Stochastic Process,” *Review of Economic Studies*, 80, 1–34.
- BÖRGERS, T. (2015): *An Introduction to the Theory of Mechanism Design*. Oxford University Press, Oxford.
- CLEMENTI, G., AND H. HOPENHAYN (2006): “A Theory of Financing Constraints and Firm Dynamics,” *Quarterly Journal of Economics*, 121, 229–265.
- COURTY, P., AND H. LI (2000): “Sequential Screening,” *Review of Economic Studies*, 67, 697–717.
- CREMER, J., AND R. MCLEAN (1985): “Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist When Demands are Interdependent,” *Econometrica*, 53, 345–361.
- (1988): “Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions,” *Econometrica*, 56, 1247–1258.
- DASGUPTA, P., AND E. MASKIN (2000): “Efficient Auctions,” *Quarterly Journal of Economics*, 115, 341–388.

- DEB, R., AND D. MISHRA (2014): “Implementation With Contingent Contracts,” *Econometrica*, 82(6), 2371–2393.
- DELLAVIGNA, S., AND U. MALMENDIER (2006): “Paying Not To Go To The Gym,” *American Economic Review*, 96(694-719).
- DEMARZO, P., AND M. FISHMAN (2007): “Optimal Long-Term Financial Contracting,” *Review of Financial Studies*, 20, 2079–2128.
- DEMARZO, P., AND Y. SANNIKOV (2006): “Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model,” *Journal of Finance*, 61, 2681–2724.
- ESŐ, P., AND B. SZENTES (2007): “Optimal Information Disclosure in Auctions,” *Review of Economic Studies*, 74, 705–731.
- (2017): “Dynamic Contracting: An Irrelevance Result,” *Theoretical Economics*, 12, 109–139.
- FARHI, E., AND I. WERNING (2007): “Inequality and Social Discounting,” *Journal of Political Economy*, 115, 365–402.
- GALLEGRO, G., AND G. VAN RYZIN (1994): “Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons,” *Management Science*, 40(8), 999–1020.
- GARRETT, D. (2016): “Intertemporal Price Discrimination: Dynamic Arrivals and Changing Values,” *American Economic Review*, 106, 3275–3299.
- GERSHKOV, A., AND B. MOLDOVANU (2009a): “Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach,” *American Economic Journal: Microeconomics*, 2, 98–168.
- (2009b): “Learning about the Future and Dynamic Efficiency,” *American Economic Review*, 99, 1576–1588.
- (2014): *Dynamic Allocation and Pricing: A Mechanism Design Approach*. MIT Press.
- GERSHKOV, A., B. MOLDOVANU, AND P. STRACK (2017): “Revenue Maximizing Mechanisms with Strategic Customer and Unknown, Markovian Demand,” *Management Science*, forthcoming.

- GOLOSOV, M., N. KOCHERLAKOTA, AND A. TSYVINSKI (2003): “Optimal Indirect and Capital Taxation,” *Review of Economic Studies*, 70, 569–587.
- GRUBB, M., AND M. OSBORNE (2015): “Cellular Service Demand: Biased Beliefs, Learning and Bill Shock,” *American Economic Review*, forthcoming.
- HARRIS, M., AND B. HOLMSTROM (1982): “A Theory of Wage Dynamics,” *Review of Economic Studies*, 49, 315–333.
- JEHIEL, P., AND B. MOLDOVANU (2001): “Efficient Design with Interdependent Valuations,” *Econometrica*, 69, 1237–1259.
- KAKADE, S., I. LOBEL, AND H. NAZERZADEH (2013): “Optimal Dynamic Mechanism Design and the Virtual Pivot Mechanism,” *Operations Research*, 61, 837–854.
- KOCHERLAKOTA, N. (1996): “Implications of Efficient Risk Sharing without Commitment,” *Review of Economic Studies*, 63, 595–609.
- KRÄHMER, D., AND R. STRAUZ (2015): “Optimal Sales Contracts with Withdrawal Rights,” *Review of Economic Studies*, 82, 762–790.
- (2017): “Sequential vs Static Screening - An Equivalence Result,” University of Bonn and Humboldt University of Berlin.
- KRASIKOV, I., AND R. LAMBA (2016): “A Theory of Dynamic Contracting with Financial Constraints,” .
- KURIKBO, N., T. LEWIS, F. LIU, AND J. SONG (2017): “Long-Term Partnership for Achieving Efficient Capacity Allocation,” .
- LI, H., AND X. SHI (2017): “Discriminatory Information Disclosure,” *American Economic Review*, forthcoming.
- LIU, H. (2013): “Efficient Dynamic Mechanisms in Environments with Interdependent Valuations“,” .
- LOBEL, I., AND R. PAES LEME (2017): “Dynamic Mechanism Design under Positive Commitment,” Discussion paper, New York University.

- MEZZETTI, C. (2004): “Mechanism Design with Interdependent Valuations: Efficiency,” *Econometrica*, 72, 1617–1626.
- (2007): “Mechanism Design with Interdependent Valuations: Surplus Extraction,” *Economic Theory*, 31, 473–488.
- MIERENDORFF, K. (2013): “The Dynamic Vickrey Auction,” *Games and Economic Behavior*, 82, 192–204.
- MILGROM, P., AND I. SEGAL (2002): “Envelope Theorem for Arbitrary Choice Sets,” *Econometrica*, 70, 583–601.
- MIRROKNI, V., R. P. LEME, P. TANG, AND S. ZUO (2017): “Non-Clairvoyant Dynamic Mechanism Design,” Discussion paper, Google Research.
- MOOKHERJEE, D., AND S. REICHELSTEIN (1992): “Dominant Strategy Implementation of Bayesian Incentive Compatible Allocation Rules,” *Journal of Economic Theory*, 56, 378–399.
- MYERSON, R. (1986): “Multistage Games with Communication,” *Econometrica*, 54, 323–358.
- PAPADIMITRIOU, C., G. PIERRAKOS, C. PSOMAS, AND A. RUBINSTEIN (2016): “The Intractability of Dynamic Mechanism Design,” in *SODA*, pp. 1458–1475.
- PAVAN, A. (2017): “Dynamic Mechanism Design: Robustness and Endogenous Types,” in *Advances in Economics and Econometrics: 11th World Congress*, ed. by M. P. B. Honore, A. Pakes, and L. Samuelson. Cambridge University Press.
- PAVAN, A., I. SEGAL, AND J. TOIKKA (2014): “Dynamic Mechanism Design: A Myersonian Approach,” *Econometrica*, 82, 601–653.
- PHELAN, C., AND R. TOWNSEND (1991): “Computing Multi-Period Information Constrained Optima,” *Review of Economic Studies*, 58, 853–881.
- RIORDAN, M., AND D. SAPPINGTON (1987): “Information, Incentives and Organizational Mode,” *Quarterly Journal of Economics*, 102, 243–264.
- ROGERSON, W. (1985): “The First-Order Approach to Principal-Agent Problems,” *Econometrica*, 51, 357–367.

SEGAL, I. (2003): “Optimal Pricing Mechanism with Unknown Demand,” *American Economic Review*, 93, 509–529.

SKRZYPACZ, A., AND J. TOIKKA (2015): “Mechanisms for Repeated Trade,” *American Economic Journal: Microeconomics*, 7, 252–293.

SPEAR, S., AND S. SRIVASTAVA (1987): “On Repeated Moral Hazard with Discounting,” *Review of Economic Studies*, 1987, 599–617.

THOMAS, J., AND T. WORRALL (1988): “Self-Enforcing Wage Contracts,” *Review of Economic Studies*, 55, 541–554.

——— (1990): “Income Fluctuation and Asymmetric Information: An Example of a Repeated Principal-Agent Problem,” *Journal of Economic Theory*, 51, 367–390.