

Dynamic Mechanism Design: An Introduction[†]

DIRK BERGEMANN AND JUUSO VÄLIMÄKI*

We provide an introduction to 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 revenue optimal mechanisms. 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 agents’ types. Third, we consider models with changing populations of agents over time. After discussing related models with risk-averse agents and limited liability, we conclude with a number of open questions and challenges that remain for the theory of dynamic mechanism design. (JEL D44, D81, D82)

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 yet precise enough to yield concrete insights and predictions. Over the last decade, the mechanism design approach has been applied to a variety of dynamic settings. In this survey, we review the basic questions and modeling issues that arise when trying to extend the static paradigm to dynamic settings. We do not aim at maximal generality of the results that we present, but we try to bring out 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 quasi-linear payoffs. Applications of this model include

* Bergemann: Yale University. Välimäki: Aalto University School of Business. We acknowledge financial support through NSF Grant 1459899 and a Google Faculty Fellowship. We would like to thank the editor, Steven Durlauf, and four anonymous referees for their extraordinarily helpful and detailed suggestions. 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 Ian Ball and Jaehee Song for valuable research assistance.

[†] Go to <https://doi.org/10.1257/jel.20180892> to visit the article page and view author disclosure statement(s).

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.

It is well-known that in dynamic principal-agent models, private information held by the agent requires the optimal contract to be a long-term arrangement, one that cannot be replicated by a sequence of short-term contracts. This is due to the “ratchet effect.” 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.

The leading example for this survey is the problem of selling a number of goods—possibly limited—over time as the demand for the goods evolves. The dynamics that arise in such problems pertain both to the evolution of the willingness to pay as well as to the set of feasible allocations over time—through a variety of natural channels:

- (i) The sales problem may be nonstationary because the goods have either a fixed supply or an expiration date like airline tickets for a particular flight date and time. The key feature here is that the opportunity cost of selling a unit of the good today is determined by the opportunities for future sales. Markets where such concerns are important, such as for airline tickets, have recently witnessed a number of new pricing practices: frequently changing prices and options for buyers to reserve a certain price for a given period of time. A vast literature under the heading *revenue management* in operations research tackles applied problems of this sort.
- (ii) Realized sales today help predict future sales if there is uncertainty

about the rate at which buyers enter the market. Professional and college sports teams base their prices for remaining tickets on the sales to date. This form of dynamic pricing is used for concert tickets, hotel booking, and transportation services, such as the surge pricing of Uber.

- (iii) The valuations of the buyers may evolve over time as they learn more about the product by using it or by observing others. Cheap trial periods for online services are a particular form of intertemporal price discrimination in this setting.
- (iv) The cost of serving the market may change over time due to exogenous improvement in technology or through learning by doing.

The general model that we consider will encompass all of these different trading environments. We cover the optimal timing of a single sale as well as repeated sales over the time horizon. In all of the above applications, the types of some agents and/or the set of allocations available change in a nontrivial manner across periods. For us, this is the distinguishing feature of dynamic mechanism design.

The techniques of dynamic mechanism design have become more prevalent in many markets over the past decade, often under the term “dynamic pricing.” The essence of dynamic pricing is to frequently adjust the price of the object over time in response to changes in the estimated demand. The optimal price is commonly adjusted through an algorithm that responds to temporal supply and demand conditions, time, competing prices, and customer attributes and behavior. The adoption of dynamic pricing strategies (particularly in e-commerce, e.g., Amazon) is facilitated by the rapid increase in real-time

data on market conditions and customer behavior, which are used to condition the price and allocation policies.

Beyond the dynamic pricing of individual items, sophisticated dynamic contracts are becoming more common. Airlines now frequently offer option contracts to allow customers to secure a certain fare for a fixed time period before they purchase the product, a “fare lock.” Many subscription services offer a trial period with a low price before the price resets at a higher level.

A different class of applications of dynamic mechanism design arises in such common situations as the pricing of memberships, such as fitness clubs, or long-term contracts, such as mobile phone contracts or equipment service contracts. At any given point in time, the potential 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 is then how to attract (and sort) the buyers with different current and future valuations for his services. The menu of possible contracts presumably has to allow the buyers to express their private current willingness to pay as well as their expectations 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 (forthcoming), 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 prices per unit of consumption; and (iii) leasing contracts where the length of the lease term is the object of choice for the consumer. We will highlight, in section 5, how these and other features of observed contract varieties may arise as solutions to dynamic mechanism design problems.

Since we insist on full commitment power throughout this survey, we bypass the vast literature on Coasian bargaining that has the lack of commitment at its heart. Similarly, we do not consider contract dynamics resulting from renegotiation. Since both the seller and the buyer commit to the mechanism, we also restrict the type of participation constraints that we allow. In particular, we do not try to give full analysis of models where the buyer’s outside option changes over time as in Harris and Holmström (1982). Dynamic games where the players engage in interactions while their payoff relevant types are subject to stochastic changes are closely related to the material in this survey.¹ Since the analysis of such games requires sequential rationality on the part of all players, our focus on the optimal commitment solutions for the designer rules them out of the scope of this survey. This survey will also not include a comprehensive review of the recent work on dynamic taxation and dynamic public finance, a line of research that has a strong focus on strictly concave, rather than quasi-linear, payoffs. However in section 7, we shall provide connections between these two classes of payoff environments and comment on the similarities and differences in the analysis and the results.

Our aim in writing this survey is twofold. We want to give an overview of the tools and techniques used in dynamic mechanism design problems in order to give the reader an understanding of the scope of applications that can be tackled within this framework. We give a number of examples where the optimal solution can be fully characterized. For many interesting dynamic contracting problems, finding optimal incentive compatible mechanisms is beyond the scope of the currently available techniques. In such cases, one must look for partial solutions or approximate

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

solutions. Second, we want to present and discuss the recent literature in this area.

We begin with mechanisms that achieve a *socially efficient* allocation. A dynamic version of the *pivot mechanism* gives each agent a private payoff equal to her marginal contribution to the utilitarian social surplus. A mechanism that has this property is attractive, since it gives each agent the societally optimal incentives to make private investments. In a dynamic context, such investments could generate more accurate valuations or reduce future costs. We give a simple formula for the periodic payments that support the efficient allocation rule. In contrast to optimal or revenue maximizing mechanisms, the dynamic pivot mechanism does not rely on strong assumptions about commitment or constant outside options.

For the case of *revenue maximizing* mechanisms, the central trade-off in the static case is between social surplus and information rent going to the agent. We investigate how far the well-known results from the static model extend. In particular, we try to stay as close as possible to the best understood model where a principal offers a contract to a privately informed agent with a single-dimensional type and supermodular preferences over allocations and types. We see conditions where the usual results of no distortions for the highest-type agent and downward distortions for all other agents hold. But sometimes the direction of the distortions may be reversed due to the type dynamics.

On the more dynamic implications, 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 type 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. We discuss a way of formalizing this line of thought and we will see the extent to which this result holds.

One of the key implications for the revenue-maximizing allocation stems from this intuition. For most stochastic processes (e.g., ergodic and strongly mixing processes), knowing the value of the process in period t tells us little about the value of the process in period $t + k$ 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 about the valuation θ_t for large t . As a result, distortions from the efficient allocation path should vanish as t becomes large. This property of the optimal allocation path is, in fact, quite robust. Whereas most of the analysis that we present relies on arguments based on the envelope theorem (i.e., arguments depending on local optimality of truthful reporting in the mechanism), the property of vanishing distortions holds also for a much wider class of models where the so-called first-order approach fails.

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 settings where efficient outcomes can be achieved. Another novel finding in this literature is that having forward-looking buyers may sometimes be good for the revenue maximization. This is very much in contrast with the typical Coasian reasoning and also represents a novel finding relative to the literature on revenue management.

In the last substantive section of this survey, we briefly consider 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 are primarily concerned with

consumption smoothing over risky outcomes. 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 quasi-linear utility. Finally, we make some connections to the rapidly growing computer science literature on mechanism design. Rather than concentrating on the properties of the optimal mechanism for a fixed stochastic model, this literature seeks mechanisms that guarantee a good payoff across a variety of different stochastic models.

The interested reader will find complementary material and more technical detail in the recent textbooks by Börgers (2015) and Gershkov and Moldovanu (2014). We have also included a few more technical observations in the appendix. The earlier survey by Bergemann and Said (2010) focuses on dynamic auctions, and the more recent survey by Pavan (2017) focuses on issues of robustness and endogenous types. Bergemann and Pavan (2015) provide an introduction to recent research in dynamic mechanism design collected in a symposium issue of the *Journal of Economic Theory*. The textbook by Talluri and Van Ryzin (2004) is a classic introduction into revenue management from the operations research perspective.

2. The Dynamic Allocation Problem

2.1 Types, Allocations, and Preferences

In this section, we present a dynamic and stochastic payoff environment that is general

enough to cover all the later sections. We consider a discounted discrete-time model with a finite or infinite ending date T . Each agent $i \in \{1, 2, \dots, I\}$ receives in each period $t \leq T$ a payoff that depends on the current physical allocation $x_t \in X_t$, the current monetary payment (or transfer) $p_{i,t} \in \mathbb{R}$, and the private information

$$\theta_t = (\theta_{i,t}, \theta_{-i,t}) \in \prod_{i=1}^I \Theta_i = \Theta \subset \mathbb{R}^I.$$

Throughout this survey, we assume private values and quasi-linear utilities. As a result, the Bernoulli utility function u_i of agent i takes the form:

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

We assume that the type $\theta_{i,t}$ of agent i follows a controlled Markov process on the state space Θ_i . The flow payoffs of the social planner are defined by:

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

The set X_t of feasible allocations in period t may depend on the vector of past allocations

$$x^{t-1} \triangleq (x_0, \dots, x_{t-1}) \in X^{t-1}.$$

For example, the seller may only have K units of the object for sale, and a sale today diminishes the number of available objects tomorrow. The dependence of the set of feasible allocations tomorrow on the current feasible set and current allocation is denoted by a transition function g :

$$X_{t+1} \triangleq g(X_t, x_t).$$

There is a common prior c.d.f. $F_{i,0}(\theta_{i,0})$ regarding the initial type $\theta_{i,0}$ of each agent i . The current type $\theta_{i,t}$ and the current action x_t determine the distribution of the type $\theta_{i,t+1}$ in the next period. We assume that

this distribution can be represented by a Markovian transition function (or stochastic kernel):

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

The utility functions u_i and the transition functions F_i are all common knowledge at $t = 0$. At the beginning of each period t , each agent i observes $\theta_{i,t}$ privately. At the end of each period, an allocation $x_t \in X_t$ is chosen by the principal 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 . To ensure that all the expectations in the model are well defined and finite, we assume that

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

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

2.2 Possible Interpretations of Types

Up to now, we have been very general about the interpretation of $\theta_{i,t}$. There are at least three separate cases that deserve mention here. In the first, all agents are present in all periods of the game, and their types evolve according to an exogenous stochastic process on Θ_i . In the second, all agents are present in all periods, but their future types depend endogenously on current allocations. In the third case, not all agents are present in all periods.

The first case seems appropriate for procuring goods over time from firms whose privately known costs follow a stochastic process $F(\theta_{i,t} | \theta_{i,t-1})$, for example, 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.

For an example of the second class of models, consider an employer, i , who learns privately about the (firm-specific) productivity

ω_i . In this case, a risk-neutral employer would compute the posterior distribution on i 's productivity:

$$\theta_{i,t} \triangleq \Pr[\omega_i | h_{i,t}],$$

where $h_{i,t}$ is the information set of firm i at time t . It makes sense to assume now that $\theta_{i,t+1}$ depends on $\theta_{i,t}$ and the allocation x_t , in particular whether the worker was employed by firm i in period t or not. Hence, the type evolution is endogenous to the allocation problem.

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 due to this more general formulation. Hence, we can accommodate other endogenous models such as learning by doing, where the production cost of a firm decreases stochastically in its past cumulative production, or habit formation and preference for variability over time.

The case where not all agents are present at all times requires a bit more discussion. An agent may, for example, decide to wait in order to get a better deal on a purchase. If all agents are present from the start, this case is covered by the previous two specifications. If agents arrive stochastically over time, they can enter contracts only after arrival. We may then assume that the arrivals are either private information to the agent or publicly observed. For the first case, we assume that each agent i can have a particular type, the null type $0 \in \Theta_i$ for all i that we interpret as indicating that the agent is not present. We assume that $v_i(x_t, 0) = 0$ for all i, x_t and that agents with type 0 cannot make or receive any transfers. The interpretation is that i is born at the first time $t = \tau$ where $\theta_{i,t} \neq 0$ and hence, his arrival time τ is private information to the agent. Alternatively, we can assume each

agent's arrival time τ is publicly observable. In section 6, we analyze and contrast such models.

2.3 Dynamic Direct Mechanisms

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 after all histories of realized and reported types. The dynamic revelation principle as first stated in Myerson (1986), and recently extended by Sugaya and Wolitzky (2017), argued that there is no loss of generality in restricting attention to dynamic direct mechanisms where the agents report their information truthfully.

The mechanism designer chooses how much of the information in the reports and allocations to disclose to the players. In this survey, we assume that the physical allocation $x_t \in X_t$ is publicly observed.² It is clear that restricting the information available to agent i makes it easier to satisfy the incentive compatibility constraints for that player. Hence, it might be easier to achieve incentive compatibility when other agents' past reports are kept secret. In all of the applications covered in this survey, we can find an optimal mechanism where all past reports and all past allocations are made public. In what follows, we consider dynamic direct mechanisms that have this feature.

Let $r_{i,t} \in \Theta_i$ denote the report of agent i in period t and let $r_t = (r_{i,t}, r_{-i,t})$ be the vector of reported types in period t . We denote the public history $(x_s, r_{i,s})_{s < t}$ at period t by $h_t \in H_t$. When agent i chooses her report in period t , she knows her own type $\theta_{i,t}$, and all her past realized types. In the Markovian setting, the only payoff relevant

private information is her current type and hence, we let the private history of agent i be $h_{i,t} = (\theta_{i,t}, h_t) \in H_{i,t}$. A reporting strategy $\mathbf{r}_i = (r_{i,t})_{t=1}^T$ of agent i is given by

$$r_{i,t} : \Theta_i \times H_t \rightarrow \Theta_i.$$

A dynamic direct mechanism $(\mathbf{x}, \mathbf{p}) = (x_t, p_t)_{t=1}^T$ assigns physical allocations and transfer payments to the agents as a function of their current reports and the public history:

$$x_t : \Theta \times H_t \rightarrow X_t,$$

$$p_t : \Theta \times H_t \rightarrow \mathbb{R}^I.$$

Notice that the reporting strategy \mathbf{r} and the allocation process \mathbf{x} induce a stochastic process for the sequence of types through the transition probability $F(\theta_{t+1} | \theta_t, x_t(r_t(\theta_t)))$. We shall be particularly interested in *truthful reporting* strategies. For this purpose, let $\hat{\mathbf{r}}$ denote the reporting strategy profile where

$$r_{i,t}(h_{i,t}) = \theta_{i,t}$$

for all $i, t, h_{i,t}$.

The physical allocation rule \mathbf{x} , and any vector \mathbf{r}_{-i} of reporting strategies for agents other than i induce a Markovian decision problem for agent i with the dynamic programming formulation:

$$\begin{aligned} & V_{i,t}(\theta_{i,t}, h_t; \mathbf{r}_{-i}) \\ &= \max_{r_{i,t}} E \left[v_i \left(x_t \left(r_{i,t}, r_{-i,t}(\tilde{\theta}_{-i,t}, h_t), h_t \right), \theta_{i,t} \right) \right. \\ & \quad \left. - p_{i,t} \left(r_{i,t}, r_{-i,t}(\tilde{\theta}_{-i,t}, h_t), h_t \right) \right. \\ & \quad \left. + \delta V_{i,t+1}(\tilde{\theta}_{i,t+1}, \tilde{h}_{t+1}; \mathbf{r}_{-i}) \right] \Big| \theta_{i,t}, h_t. \end{aligned}$$

²In some cases, such as allocating a fixed capacity over time, each player is only informed of her own allocation and the number of remaining units may be kept secret.

We designate random variables, such as $\tilde{\theta}_{i,t}$, by tilde. The expectation above is taken with respect to the stochastic process $\{\tilde{\theta}_t\}$ induced by the transition probability F , the allocation rule x , and the reporting strategy \mathbf{r} . For the remainder of the paper we shall suppress the time index when the conditioning variable implicitly defines the time index for the relevant function. Thus, for example, we will write $V_i(\theta_{i,t}, h_t; \mathbf{r}_{-i})$ rather than $V_{i,t}(\theta_{i,t}, h_t; \mathbf{r}_{-i})$, and $x(r_{i,t}, r_{-i,t}(\tilde{\theta}_{-i,t}, h_t), h_t)$ rather than $x_t(r_{i,t}, r_{-i,t}(\tilde{\theta}_{-i,t}, h_t), h_t)$.

We define

$$V_{i,t}(\theta_{i,t}, h_t) \triangleq V_{i,t}(\theta_{i,t}, h_t; \hat{\mathbf{r}}_{-i})$$

to be the value function of agent i under truthful reporting by the other agents. We say that (\mathbf{x}, \mathbf{p}) is Bayes–Nash implementable if for all $i, t, h_{i,t}$,

$$\begin{aligned} \theta_{i,t} \in \arg \max_{r_{i,t}} E \Big[& v_i(x(r_{i,t}, \tilde{\theta}_{-i,t}, h_t), \theta_{i,t}) \\ & - p_i(r_{i,t}, \tilde{\theta}_{-i,t}, h_t) \\ & + \delta V_i(\tilde{\theta}_{i,t+1}, \tilde{h}_{t+1}) \Big] \Big| \theta_{i,t}, h_t \Big]. \end{aligned}$$

Thus, taking the expectation over the other agents' true type realization truth telling in every period and after every history is an optimal strategy for agent i .

We will sometimes refer to a stronger notion of implementability, called *periodic ex post* implementability. To define this notion, we let

$$\begin{aligned} V_{i,t}(\theta_t, h_t) \\ \triangleq \max_{r_{i,t}} \Big\{ & v_i(x(r_{i,t}, \theta_{-i,t}, h_t), \theta_{i,t}) - p_i(r_{i,t}, \theta_{-i,t}, h_t) \\ & + \delta E \left[V_i(\tilde{\theta}_{t+1}, \tilde{h}_{t+1}) \Big| \theta_t, h_t \right] \Big\}. \end{aligned}$$

Observe that the value function is now defined on the set of type *vectors*, θ_t , and public histories, rather than just the individual type $\theta_{i,t}$. A mechanism (\mathbf{x}, \mathbf{p}) is periodic ex post implementable if for all i, t, θ_t, h_t :

$$\begin{aligned} \theta_{i,t} \in \arg \max_{r_{i,t}} \Big\{ & v_i(x(r_{i,t}, \theta_{-i,t}, h_t), \theta_{i,t}) \\ & - p_i(r_{i,t}, \theta_{-i,t}, h_t) \\ & + \delta E \left[V_i(\tilde{\theta}_{t+1}, \tilde{h}_{t+1}) \Big| \theta_t, h_t \right] \Big\}. \end{aligned}$$

Whenever (\mathbf{x}, \mathbf{p}) is periodic ex post implementable, no agent wants to change her report after learning the contemporaneous reports of the other agents. This means that as in the static setting, ex post implementation is a solution concept that is stronger than Bayes–Nash implementation but weaker than dominant strategy implementability.

3. Efficient Dynamic Mechanisms

We begin the analysis by describing three dynamic mechanisms that attain the intertemporally efficient allocation in the presence of private information arriving over time: (i) the team mechanism, (ii) the dynamic pivot mechanism, and (iii) the dynamic AGV mechanism. We illustrate these mechanisms by considering allocating a fixed number of objects over time.

3.1 The Team Mechanism

We start by constructing a simple mechanism that makes truthful reporting incentive compatible in the sense of periodic ex post incentive compatibility. In this mechanism, called the team mechanism, the agents have the right incentives to report their types truthfully as their payoff is the entire social surplus generated in the allocation problem. Hence, the right place to start the construction is the social planner's optimal allocation problem with publicly observable types. The

utilitarian welfare maximization, including all agents and the principal, solves the following program:

$$W(\theta_0, X_0) \triangleq \max_{\{x_t \in X_t\}_{t=0}^T} E \left\{ \sum_{t=0}^T \delta^t \sum_{i=0}^I v_i(x_t, \theta_{i,t}) \right\},$$

where the expectation is taken with respect to $F(\theta_{t+1} | \theta_t, x_t)$ and the feasibility condition $x_t \in X_t$ for all t . Notice, that we allow the value function in period 0 (and in all future periods) to depend explicitly on the set of feasible allocations. We define the social flow payoff as:

$$w(x_t, \theta_t) \triangleq \sum_{i=0}^I v_i(x_t, \theta_{i,t}).$$

We can write this in terms of a dynamic program:

$$W(\theta_t, X_t) = \max_{x_t \in X_t} \{w(x_t, \theta_t) + \delta E W(\theta_{t+1}, X_{t+1})\}$$

subject to the transition function of the state

$$\theta_{t+1} \sim F(\cdot | \theta_t, x_t),$$

and the feasibility constraint:

$$X_{t+1} = g(X_t, x_t).$$

Let $x_i^*(\theta_t, X_t)$ denote an optimal policy for this program.

As in the static setting, periodic ex post incentive compatibility follows if we give each agent the entire social surplus. By the one-shot deviation principle, it is sufficient to set for all i and all θ_i :

$$\begin{aligned} p_i^*(r_{i,t}, r_{-i,t}) &= - \sum_{j \neq i} v_j(x^*(r_{i,t}, r_{-i,t}), r_{j,t}) \\ &\triangleq -w_{-i,t}(r_t). \end{aligned}$$

In other words, each agent is paid in each period the efficient gross surplus that the other agents receive at the efficient allocation.

Up to now, we have allowed correlated types, and in fact, $(\mathbf{x}^*, \mathbf{p}^*)$ is periodic ex post incentive compatible for correlated as well as for independent types. Strengthening the notion of incentive compatibility to dominant strategies is, unfortunately, not possible. The easiest way to see this is to notice that opponents' future reports depend on past allocations (and possibly also on other agents' past reports). To ensure the implementability of efficient allocations in future periods, a Vickrey–Clark–Groves (VCG) term, $w_{-i}(r_s)$, must be included in all future transfers.

We emphasize that we only ensure periodic ex post incentive compatibility, but not dominant strategy incentive compatibility. We illustrate by means of an example in section 3.3 that the notion of dominant strategy incentive compatibility is typically too demanding in dynamic settings.

For periodic ex post incentive compatibility, we can also allow for interdependent types between the players as long as the payoff consequences to i resulting from information of player j become observable at some later point in time.³ Unfortunately, the team mechanism results in a deficit of size $(I - 1)W_0(\theta_0, X_0)$. In the next two subsections, we consider efficient mechanisms that reduce and sometimes eliminate this deficit. We shall see that modified versions of the pivot mechanism and the and D'Aspremont–Gérard–Varet (AGV) mechanism in the static setting perform well within the dynamic model.

3.2 *Leading Example: Sequential Allocation of Fixed Capacity*

Throughout this survey, we illustrate the results with the following sequential allocation problem. There is a fixed supply of K identical indivisible goods at $t = 0$

³In this way, the team mechanism can be extended to cover the mechanisms first displayed in Mezzetti (2004).

and there is uncertainty about the demand, which is realized stochastically over time.

There are I potential bidders, each with unit demand, over two periods, $t \in \{1, 2\}$. For each $i \in \mathcal{I} \triangleq \{1, \dots, I\}$, we let $x_{i,t} = 1$ denote the event that a good is allocated to agent i in period t , that otherwise $x_{i,t} = 0$. The allocation is once and for all: after one of the goods has been allocated to i , it cannot be taken away and allocated to another bidder. By x_t , we denote the vector of allocation decisions in period t . The capacity constraint states that $\sum_t \sum_i x_{it} \leq K$.

At the beginning of period 1, each bidder observes his type $\theta_{i,1} \in [0, 1]$. The type $\theta_{i,2} \in [0, 1]$ in period $t = 2$ depends on the vector of realized types in period 1, θ_1 . In a dynamic direct mechanism, the agents report their realized types in each period. The allocation rule depends on both of these announcements. The payoff of agent i in the mechanism is $\sum_t \delta^{t-1} (\theta_{i,t} x_{i,t} - p_{i,t})$.

We denote the number of unallocated goods in period $t = 2$ by $K_2 = K - \sum_i x_{i,1}$. Efficient allocation in $t = 2$ requires allocating the K_2 units to those bidders that have the highest $\theta_{i,2}$ amongst the ones with $x_{i,1} = 0$. Solving for the efficient allocation in $t = 1$ is not trivial even in this simple allocation problem. Let \mathcal{I}_1 denote the set of bidders that receive the good in period 1 and let K_1 denote the number of goods allocated in $t = 1$. Let $\mathcal{I}_2 = \mathcal{I} \setminus \mathcal{I}_1$, and we can write the value function for the efficient period 2 allocation as:

$$W_2(\theta_2, \mathcal{I}_2, K_2) = \max_{\{x_{i,2}\}} \sum_{i \in \mathcal{I}_2} \theta_{i,2} x_{i,2},$$

subject to $\sum_i x_{i,2} \leq K_2$. Similarly, for the first period, we can write

$$W_1(\theta_1, \mathcal{I}, K) = \max_{\{x_{i,1}\}} E \left[\sum_{i \in \mathcal{I}} \theta_{i,1} x_{i,1} + \delta W_2(\mathcal{I}_2, K - K_1) \mid \theta_1 \right]$$

subject to $\sum_i x_{i,1} = K_1$.

The first-period decision fixes first the number K_1 to be allocated in $t = 1$ and then asks how to optimally choose the identities of the K_1 agents to receive the goods in $t = 1$. The next step optimizes over K_1 . The hard step is obviously in determining the set of agents to receive the goods in the first period. Unless we specify the model further, little can be said in general about the features of the optimal allocation decision. The first-period decision incorporates a few dynamic considerations. Bidder i may be present only in $t = 1$, in which case $\theta_{i,2} = 0$, or alternatively she could only arrive in $t = 2$, in which case we could take $\theta_{i,1} = 0$. Bidder i 's true valuation may be learned in $t = 2$, in which case, we could set $\theta_{i,2} \sim F_i(\cdot \mid \theta_{i,1})$. Nevertheless, since the allocation choices are finite in the above problem, an optimal allocation policy \mathbf{x}^* exists and it is easy to see that by setting

$$p_{i,t}^*(r_{i,t}, r_{-i,t}) \triangleq -w_{-i}(r_t),$$

the mechanism $(\mathbf{x}^*, \mathbf{p}^*)$ is periodic ex post incentive compatible.

3.3 Impossibility of Dominant Strategy Implementation

We first illustrate the impossibility of implementing the efficient allocation rule in dominant strategies by specifying the above example even further.

We consider two bidders, thus $I = 2$, who compete for a single object, thus $K = 1$, in a two-period model. We suppose further that bidder 1 draws a valuation θ_1 uniformly from $[0, 1]$ in period 1. Her valuation for the good remains unchanged (but is discounted) in period 2. Bidder 2 is active only in period 2, in the sense that her valuation for allocation in $t = 1$ is 0 with probability 1. Her valuation θ_2 in period $t = 2$ is independent of the valuation of bidder 1 and also drawn from the uniform distribution on $[0, 1]$. The

allocation decision is nontrivial in $t = 2$ only if the good was not allocated in $t = 1$, i.e., if $K_2 = 1$. In this case, $x_{1,2}^*(r) = 1$ and $x_{2,2}^*(r) = 0$ if $r_{1,1} \geq r_{2,2}$, otherwise $x_{1,2}^*(r) = 0$ and $x_{2,2}^*(r) = 1$.

The efficient allocation in $t = 1$ gives:

$$x^*(\theta_1) = 1 \Leftrightarrow \theta_1 \geq \delta E\left[\max\{\theta_1, \tilde{\theta}_2\}\right],$$

which results in a threshold value θ^* to allocate the object to the first bidder in the first period if:

$$\theta_1 > \theta^* = \frac{1}{\delta} - \sqrt{\frac{1}{\delta^2} - 1}.$$

The threshold θ^* is strictly increasing in the discount factor δ as the future realization of θ_2 provides a greater option value with a higher discount factor.

Green and Laffont (1977) and Holmström (1979) show that the transfer rules for efficient mechanisms are uniquely pinned down (up to a constant) when the type sets are path-connected. Combining this with the logic of VCG mechanisms allows us to conclude that in $t = 1$, the expected transfer of bidder 1 depends on her reported type $r_{1,1}$ only through its impact on the efficient allocation:

$$(1) \quad E\left[p_1(r_{1,1}, r_{2,2})\right] \\ = -E\left[\delta r_{2,2} x_{2,2}^*(r_{1,1}, r_{2,2}) + \phi_1(r_{2,2})\right]$$

where the (arbitrary) component function $\phi_1(r_{2,2})$ that enters the price p_1 depends only on the report of bidder 2.

The payment p_1 of bidder 1 displays no interaction between her report and the report of bidder 2 beyond their joint effect on the payoff to bidder 2 in the efficient allocation rule, the first term on the right-hand side of (1). Now, in order to secure the efficient decision in period $t = 1$, the transfer in equation (1) must be calculated using the truthful reporting strategy for bidder 2 (and hence the true distribution of θ_2) in $t = 2$.

But now we can show that these transfers can't possibly implement truthful reporting as a dominant strategy for player 1. To see this, suppose that bidder 1 expects that bidder 2 does not report truthfully, but rather reports $r_{2,2} = 0$ for all θ_2 , and thus

$$E\left[r_{2,2} x_{2,2}^*(r_{1,1}, r_{2,2})\right] = 0.$$

Now truthful reporting does not constitute an optimal report in period $t = 1$. Under this candidate reporting strategy of bidder 2 and given the payment rule, bidder 1 would optimally exaggerate her type in period $t = 1$ and report $r_{1,1} > \theta^*$ for all valuation $\theta_1 > 0$ since she would prefer an early allocation of the object, or

$$\theta_1 > \delta \theta_1,$$

and given her expectation about the reporting strategy of bidder 2, she would not forego any compensation she might have received if the object were to be allocated to bidder 2 in period 2. Thus, dominant strategy implementation is impossible to guarantee even in this elementary two period allocation problem.

Notice that a similar impossibility argument would emerge if bidder 1 would expect bidder 2 to report $r_{2,2} = 1$ for all θ_2 . In this case, bidder 1 would not report truthfully, but rather downward misreport in order to obtain the compensatory payments due to the high report of bidder 2 since now bidder 1 would expect to get $E\left[r_{2,2} x_{2,2}^*(r_{1,1}, r_{2,2})\right] = 1$.

Thus, in dynamic mechanism design problems, incentive compatibility in dominant strategies is too demanding. This example also shows that implementing the efficient allocation rule will typically not be detail free, and in particular will depend on the common prior regarding the distribution of the types. In the current example, the efficient allocation decisions in the initial

period (the computation of θ^*) depended on the distribution of future types, and as a result, the optimal mechanism—in particular the transfer function—also reflects the informational details of the valuation process.⁴

3.4 The Dynamic Pivot Mechanism

For the remainder of this section, we concentrate on two particular efficient mechanisms that have further desirable properties. We begin with the dynamic pivot mechanism, introduced in Bergemann and Välimäki (2010), which ensures that each agent's payoff in the mechanism corresponds to her marginal contribution to the societal welfare 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. In the next subsection, we consider the dynamic counterpart of the AGV mechanism where the focus shifts toward 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.

We now construct the dynamic pivot mechanism for the general model described in section 2 under the assumption of independent private values. We give an example in the appendix showing that dynamic pivot mechanisms do not always exist if the values are correlated. We recall that in the static pivot mechanism—introduced by Green and Laffont

(1977)—the transfers are constructed as follows:

$$(2) \quad p_i(\theta) = -\sum_{j \neq i} v_j(x^*(\theta), \theta_j) + \sum_{j \neq i} v_j(x_{-i}^*(\theta_{-i}), \theta_j),$$

where $x_{-i}^*(\theta_{-i})$ is the optimal allocation if agent i is not participating, thus

$$x_{-i}^*(\theta_{-i}) \in \arg \max_{x_t \in X_t} \sum_{j \neq i} v_j(x_t, \theta_{j,t}).$$

The idea behind the pivotal transfers is to equate each agent's expected payoff to her expected contribution to the social value. At state (θ_t, X_t) , we compute the (dynamic) marginal contribution of agent i :

$$M_i(\theta_t, X_t) \triangleq W(\theta_t, X_t) - W_{-i}(\theta_t, X_t),$$

where W and W_{-i} are the value functions of social surpluses with and without i in the society (in all future periods), respectively.

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 $m_i(\theta_t, X_t)$ in each period. The flow marginal contribution accrues incrementally over time and is defined recursively:

$$(3) \quad M_i(\theta_t, X_t) = m_i(\theta_t, X_t) + \delta E M_i(\theta_{t+1}, X_{t+1}).$$

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

$$(4) \quad p_{i,t}^*(\theta_t, X_t) \triangleq v_i(x_t^*, \theta_{i,t}) - m_{i,t}(\theta_t, X_t).$$

⁴The failure of dominant strategy implementability in dynamic mechanism is similar to the failure of the static AGV mechanism to be dominant strategy incentive compatible. In both instances, it is critical to use the same expectations regarding the behavior of the other agents in the determination of the efficient rule and in the computation of expected payments for an individual agent.

We refer to $p_i^*(\theta_t, X_t)$ as the transfer of the dynamic pivot mechanism. Notice that in contrast to the static transfer payment, the reported type of agent i has also an indirect effect through $\delta EW_{-i}(\theta_{t+1}, X_{t+1})$.⁵ 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 $W(\theta_t, X_t) - W_{-i}(\theta_t, X_t)$, 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.*

It should be noted that as in any dynamic context, it is hard to pin down the exact timing of payments. Making a payment of p in period t has the same payoff consequences as making a payment of p/δ in $t+1$. In Bergemann and Välimäki (2010), we give sufficient conditions for the uniqueness of the above payment rule. Similar to the static case (Moulin 1986), these conditions require a rich domain of possible preferences. In the dynamic context, this also requires an assumption that amounts to allowing the agents to leave the mechanism stochastically. By this we mean that after leaving, no more transfers can be enacted.

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

⁵Since $W_{-i}(\theta_t, X_t) = \max_{x_t \in X_t} \{w_{-i}(x_t, \theta_t) + \delta EW_{-i}(\theta_{t+1}, X_{t+1})\}$.

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 illustrate how the payments in the dynamic pivot mechanism are computed for the leading example of section 3.3.

EXAMPLE 1 (Dynamic Pivot Mechanism for Fixed Capacity Allocation):

We first compute the marginal contributions of the agents. If agent 1 is not present, then the expected social surplus is the discounted expected value of the good to bidder 2, i.e., $\delta/2$. The expected social surplus at any moment in time when bidder 2 is not present is simply θ_1 . If the good has already been allocated, then there is no social surplus for the continuation problem. Without loss of generality, we can restrict attention to integer allocations, $x \in \{0, 1\}$. The marginal contributions of bidder 1 in periods 1 and 2 are then:

$$M_1(\theta_1, 1) = \begin{cases} \theta_1 - \delta/2, & \text{if } x_{1,1}^*(\theta_1) = 1, \\ \delta E \max\{\theta_1 - \theta_2, 0\}, & \text{otherwise;} \end{cases}$$

$$M_1(\theta_1, \theta_2, 1) = \max\{\theta_1 - \theta_2, 0\};$$

and similarly for bidder 2:

$$M_2(\theta_1, 1) = \begin{cases} 0, & \text{if } x_{1,1}^*(\theta_1) = 1, \\ \delta E \max\{\theta_2 - \theta_1, 0\}, & \text{otherwise;} \end{cases}$$

and

$$M_2(\theta_1, \theta_2, 1) = \max\{\theta_2 - \theta_1, 0\}.$$
⁶

⁶We recall that, earlier, we adopted the convention that we suppress the time index when the conditioning variable implicitly defines the time. Thus, $M_1(\theta_{1,1})$ is the marginal contribution of player 1 if the object has not been allocated yet, i.e., $= K_{11}$, in period 1, since the only conditioning variable used is the type of bidder 1, which is known in period 1.

Using the implicit definition of the flow marginal contribution given by (3), we get

$$m_1(\theta_1, 1) = \begin{cases} \theta_1 - \delta/2, & \text{if } x_1^*(\theta_1) = 1, \\ 0, & \text{otherwise.} \end{cases}$$

By equating

$$p_1(\theta_1, 1) = v_1(x_{i,t}^*, \theta_{i,t}) - m_{i,1}(\theta_{1,1}, 1),$$

we get

$$p_1(\theta_1, 1) = \begin{cases} \delta/2, & \text{if } x_1^*(\theta_1) = 1, \\ 0, & \text{otherwise.} \end{cases}$$

Similarly,

$$p_1(\theta_1, \theta_2, 1) = \theta_2, \quad \text{if } \theta_1 \geq \theta_2,$$

$$p_1(\theta_1, \theta_2, 1) = 0, \quad \text{otherwise;}$$

$$p_2(\theta_1, \theta_2, 1) = \theta_1, \quad \text{if } \theta_2 > \theta_1,$$

$$p_2(\theta_1, \theta_2, 1) = 0, \quad \text{otherwise.}$$

This suggests a simple indirect implementation of the efficient two-period auction. Bidder 1 is given the option of purchasing the good at the opportunity cost $\delta/2$ of allocating the good in $t = 2$. If she does not exercise the option, then the good is sold in a second price auction without reserve prices in period 2. It should be noted that finding the right price for this indirect implementation is remarkably easy in comparison to finding the efficient threshold type θ^* .

More generally, it is quite easy to compute the direct version of the dynamic pivot mechanism on the basis of the dynamic social surpluses using dynamic programming techniques. For example, Bergemann and Välimäki (2003, 2006) use the construction of the dynamic marginal contribution to solve for the equilibrium of dynamic common

agency problems and dynamic competition problems, respectively. These earlier contributions considered symmetric but imperfect information environments; Bergemann and Välimäki (2010) establish that the underlying principles extend to asymmetric information environments as well, and then can solve priority queuing problems as analyzed by Dolan (1978).

Since we have assumed independent types, additional assumptions on the connectedness of the type spaces and payoff functions guarantee a dynamic payoff equivalence result via the envelope theorem of Milgrom and Segal (2002). By imposing an individual rationality or participation constraint for the agents, it is often possible to show that similar to the static setting, the dynamic pivot mechanism maximizes the expected transfers from the agents 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 individual rationality. Skrzypacz and Toikka (2015) consider a model of repeated bilateral monopoly with varying degrees of persistence for the buyer's valuations and the seller's costs. With perfectly persistent types, the Myerson–Satterthwaite theorem applies and efficient trade is impossible. With independent types, the expected gains from future trades can be used to relax the participation constraint and efficient trading may become possible. Different levels of persistence then determine different sets of efficient budget balanced trading rules for the problem.

3.5 Balancing the Budget

The static VCG mechanism is defined by $(x^*(\theta), p(\theta))$, and for all i and all θ , we have

$$p_i(r) = -w_{-i}(r) + \phi_i(r_{-i}),$$

where the second component of the transfer function, $\phi_i(r_{-i})$ of agent i is an arbitrary function that 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. The static VCG mechanisms are dominant strategy incentive compatible, i.e., they induce truth telling as a best response against any reported type vector of other agents. As a result, a modified mechanism where

$$p_i(r) = -E_{\theta_{-i}}[w_{-i}(r_i, \theta_{-i}) + \phi_i(\theta_{-i})],$$

and the expectation is taken with respect the marginal distribution on the other agents' types (recall that we have assumed independence here) is incentive compatible as long as the other agents announce their types truthfully. Budget balance is obtained by specifying (with the understanding that $I + 1 = 1$):

$$\phi_{i+1}(r_{-(i+1)}) \triangleq E_{r_{-i}}[w_{-i}(x^*(r_i, r_{-i}), r_{-i})].$$

Notice that we have to give up on dominant strategy incentive compatibility here. If the other players lie about their type, the first term in the transfer does not equate the bidder's payments to the social surplus and hence it may well be that lying is optimal for i too. Observe the similarity of this reasoning to the failure of dominant strategy incentive compatibility in the dynamic example above.

For the dynamic mechanism, the transfer payments must be constructed in such a way that similar problems do not arise due to the dynamic nature of the announcements. Supposing that the incentive payments are made as above, based on the expectations over other players' types, but that the realizations of other players' types can be inferred from the allocations prior to one's own announcement, the simple

AGV-mechanism is no longer incentive compatible. In order to secure incentive compatibility, Athey and Segal (2013) modify the transfers to overcome this problem by aligning agent i 's incentive pay with the change in the expected externality on the other agents resulting from i 's report.

The resulting balanced budget mechanisms can be quite complicated and it may not be easy to find natural indirect mechanisms for their implementation. One instance where this can be done is in the context of dynamically allocating the capacity shares in a joint project when private information about future profits arrives over time. In a dynamic sharing problem, Kuribko et al. (2017) find a version of the dynamic mechanism with budget balance that can also handle individual rationality constraints.

3.6 Interdependent Values and Correlation

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. To see this, consider the following simple version of the famous lemons problem with common values.

EXAMPLE 2 (Common Values versus Correlated Private Values):

Two agents decide the allocation of an indivisible object in a two-period model. The allocation $x_t \in \{1, 2\}$ for $t \in \{1, 2\}$ records who gets the object in period t . Agent 1 is privately informed of the quality of the object, i.e., $\theta_1 \in [0, 2]$, and her value from consuming the good in $t = 1$ is given by her private type θ_1 . If the good is allocated to agent 2 in $t = 1$, then agent 2 receives her gross utility $\theta_2 = 2\theta_1 - 1$ after consuming the good in period $t = 1$. In $t = 2$, agent 2 just reports her type θ_2 and transfers are made.

If all transactions and transfers take place in $t = 1$, then the efficient allocation rule

where agent 1 consumes the good if and only if $\theta_1 \leq 1$ is not monotone and hence not incentive compatible in the static mechanism. By allowing agent 2 to learn her value after consuming the good, the common values lemons model becomes a dynamic private values model since, conditional on knowing θ_2 , player 2 does not care about θ_1 . As a result, we can compute the efficient team mechanism to support trade here. The reported θ_1 determines the allocation: there is trade if and only if $\theta_1 \geq 1$ and the reported θ_2 determines the payment $t_1 = \theta_2$ that agent 1 receives in $t = 1$.

The idea that incentive constraints can be relaxed by using future realizations of correlated signals is elaborated further in Mezzetti (2004, 2007) and Deb and Mishra (2014). For many applications, experienced utilities after trade are natural signals of this type. It should also be noted that there are often quite natural implementations of the efficient mechanisms. Contracts with money-back clauses can be used to facilitate trade that would otherwise be limited by the lemons problem.

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 Crémer and McLean (1985, 1988). Liu (2013) also covers the case of interdependent values but independently distributed signals and develops a dynamic version of the generalized VCG mechanism along the lines of the static version of Dasgupta and Maskin (2000).

4. *Optimal Dynamic Mechanisms*

We now shift our attention from socially efficient mechanisms to revenue-maximizing, or (revenue) optimal mechanisms. In static

environments, the key economic insight is the resulting trade-off between efficiency and information rent left to the agents. In socially efficient mechanisms, this trade-off is absent due to quasi-linear preferences. After all, the utilitarian solution does not preclude information rents. By contrast, if a seller tries to maximize her sales revenue or if a regulator puts a lower weight on profits than on consumer surplus in her objective function, then this trade-off emerges. Deviations from the surplus-maximizing allocations are generally optimal, since the reductions in information rent to the privately informed parties more than compensate for the losses in the social surplus. For dynamic models of mechanism design, the key issue is then how information rent accrues to privately informed agents over the contracting horizon.

Consider a model where a privately informed agent contracts with an uninformed principal at the beginning of a dynamic allocation problem. If the agent knows all her future information types, the model is a multidimensional mechanism design problem and as it is well-known, it is very hard to characterize the optimal contract in such environments. Fortunately, it is often quite reasonable that the agent does not know her future types. Of course, her current type allows her to predict her future types more accurately than the principal (except in the less interesting case of i.i.d. types). The main analytical challenge in optimal dynamic contracting problems is to characterize how the initial private information affects the future information rents and how the optimal allocation trades off these effects relative to the social surplus maximizing allocation.

Optimal dynamic contracting 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 a privately known marginal cost in the first period and where the second-period marginal cost is unknown to both parties (but may depend stochastically on the first period cost). Within a similar model, Riordan and Sappington (1987) analyze task assignment in a similar two-period principal–agent model. Besanko (1985) covers a dynamic regulation problem where the privately known cost parameter is either i.i.d. over time or follows a first-order autoregressive process.

In the past decade, the literature has developed considerably beyond these early contributions. Much of the early literature was focused on the case where the allocation problem itself is assumed to be time invariant in the sense that the set of feasible choices in t is independent of allocation decisions for $s < t$. These papers also assumed that the distribution of future types is independent of current allocation decisions. The first assumption is violated in any dynamic problem of capacity allocation, and also in models with a fixed decision date but a dynamic flow of private information prior to the decision. The second assumption is violated in models with endogenous learning about the payoffs from a fixed set of alternatives. Examples of this type include dynamic assignment of workers to tasks or dynamic sales of experience goods.

The first analysis of the revenue-maximizing dynamic sales problem with a fixed selling date appears in Courty and Li (2000) under the name of *sequential screening*. Board (2007) extended the analysis to the case where the sales date itself is also endogenously chosen. Battaglini (2005) considers a nonlinear pricing model (with variable quantity or quality) in which the buyer's valuation changes 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) present a general infinite-horizon model that allows for general allocation problems and endogenous type processes. Their model encompasses the earlier literature with continuous type spaces and emphasizes the connections to static allocation problems of the Myersonian type. They obtain necessary conditions for incentive compatibility and present a variety of sufficient conditions for revenue-maximizing contracts for specific classes of environments.

Our goal here is to use the tools from static mechanism design as much as possible to understand the basic analytics of the dynamic problem. Hence we start with a brief review of the main results in static mechanism design. Then we connect the static and dynamic formulations by transforming the original dynamic problem into an equivalent one where it is easier to separate the initial private information and future information that can be taken to be independent of the initial type, as proposed in Esó and Szentes (2007). After formulating the optimization problem of the principal, we discuss some examples with explicit solutions to the optimal contracting problem.

4.1 *Preliminaries from Static Mechanism Design*

Not surprisingly, the assumptions needed for tractability in static mechanism design problems are also needed in the dynamic case. We assume that all payoff functions are linear in transfers, the agents' types are drawn from intervals of the real line, $\Theta_i = [\underline{\theta}_i, \bar{\theta}_i] \subset \mathbb{R}$, $\tilde{\theta}_i$ is independent of $\tilde{\theta}_j$, and that the agents' payoff functions $v_i(x, \theta)$ are strictly supermodular in (x, θ_i) . To make the connection to the dynamic setting more immediate, we allow the allocation decision $x \in X$ to be multidimensional. In the presentation below, we concentrate on the case

with a single agent and therefore we omit the subscripts.⁷ 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

$$\begin{aligned} U(\theta; \theta) &\triangleq v(x(\theta), \theta) - p(\theta) \\ &\geq v(x(r), \theta) - p(r) \triangleq U(\theta; r), \end{aligned}$$

and let

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

The envelope theorem by Milgrom and Segal (2002) gives the following necessary condition for incentive compatibility, their theorems 2 and 3, respectively. We denote by $v_2(x, \theta)$ the partial derivative with respect to the second argument, here θ .

THEOREM 2 (Payoff Equivalence): *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 ,*

$$|v_2(x, \theta)| \leq K.$$

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

$$(5) \quad V(\theta) = V(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} v_2(x(s), s) ds.$$

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

$$(6) \quad \begin{aligned} p(\theta) &= v(x(\theta), \theta) - V(\underline{\theta}) \\ &\quad - \int_{\underline{\theta}}^{\theta} v_2(x(s), s) ds. \end{aligned}$$

⁷Since types are independent, incentive compatibility reduces to individual incentive compatibility by taking expectations over the other agents' types.

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

$$(7) \quad \int_r^{\theta} (v_2(x(s), s) - v_2(x(r), s)) ds \geq 0.$$

As an intermediate step toward the dynamic analysis, consider for a moment the case where the principal and the agent can write contracts based on a publicly observable sequence of random variables, $\{\tilde{\varepsilon}_t\}_{t=1}^T$, so that the allocation x_0 is determined by the initial private information θ and x_t is determined by θ and the sequence of realizations $\varepsilon^t \triangleq (\varepsilon_1, \dots, \varepsilon_t)$. Assume also that the following separability requirement is satisfied:

$$v_t(x, \theta, \varepsilon^T) = v_t(x_t, \theta, \varepsilon^t).$$

Incentive compatibility is then equivalent to requiring that for all $\theta, r \in \Theta$,

$$(8) \quad \int_r^{\theta} E_{\varepsilon} \left[\sum_{t=0}^{\infty} \delta^t \frac{\partial v_t}{\partial \theta}(x_t(s, \tilde{\varepsilon}^t), s) - \sum_{t=0}^{\infty} \delta^t \frac{\partial v_t}{\partial \theta}(x_t(s, \tilde{\varepsilon}^t), r) \right] ds \geq 0.$$

In the static setting, single-dimensional types and single-dimensional allocations with supermodular payoff functions yield a full characterization of incentive compatible allocation rules: a mechanism is incentive compatible if and only if the physical allocation is monotone and the transfers are pinned down by the payoff equivalence theorem. The static problem is then solved by maximizing the designer's objective function over the feasible set of monotone allocation rules.

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. The static formula for contingent allocations in

equation (8) requires monotonicity on average when taking expectations over $\{\tilde{\varepsilon}_t\}_{t=1}^T$. If $x_t(\theta, \tilde{\varepsilon}^t)$ is monotonic in θ for all realizations ε^t , then incentive compatibility follows. This sufficient condition is obviously not a necessary condition, since the sum of non-monotonic functions may well be monotonic.

In the remainder of this section, we express the process of types θ_t in terms of the initial type θ_0 and a sequence of independent (uniform) random variables $\{\tilde{\varepsilon}_t\}_{t=1}^T$. The main analysis considers the case where the $\{\tilde{\varepsilon}_t\}_{t=1}^T$ are privately observed by the agent, but we also discuss the connections to the static case, above, where all future information $\{\tilde{\varepsilon}_t\}_{t=1}^T$ is publicly observed.

4.2 Orthogonalized Information

Esó and Szentes (2007) emphasize the benefits from distinguishing between private information at the time of contracting (captured in θ_0) and subsequent independent private information. With Markovian types, each θ_t is statistically dependent on θ_{t-1} and as a result, all future types are also influenced by the initial type θ_0 . We present, below, an equivalent formulation for the type process $\{\tilde{\theta}_t\}_{t=0}^T$ that allows us to distinguish in a transparent manner between the initial and the future private information despite the statistical dependency.

Consider an arbitrary random variable \tilde{x} with distribution function F . Then the random variable \tilde{y} , with

$$\tilde{y} \triangleq F(\tilde{x}),$$

is by construction uniformly distributed on $[0, 1]$. Building on this simple observation, consider next a random variable $\tilde{\theta}_1$ with a conditional distribution $F(\cdot|\theta_0)$ dependent on some realization θ_0 . Then

$$\tilde{\varepsilon}_1 \triangleq F(\tilde{\theta}_1|\theta_0)$$

is uniformly distributed for all θ_0 . As a result, $\tilde{\varepsilon}_1$ is independent of θ_0 by construction. We can view ε_1 as the realized percentile in the conditional distribution $F(\cdot|\theta_0)$. Since $F(\cdot|\theta_0)$ is an increasing function, knowledge of ε_1 and θ_0 allows for solving $\theta_1 = F^{-1}(\varepsilon_1|\theta_0)$.⁸ Hence, the information content of $(\tilde{\theta}_0, \tilde{\theta}_1)$ is the same as the information content of $(\tilde{\theta}_0, \tilde{\varepsilon}_1)$.

As a final preparatory step, let us consider the effect of $\tilde{\theta}_0$ on $\tilde{\theta}_1$. Letting $\theta_1 = F^{-1}(\varepsilon_1|\theta_0)$, we can evaluate the effect of a change in θ_0 on θ_1 for a fixed ε_1 :⁹

$$\frac{\partial F^{-1}(\varepsilon_1|\theta_0)}{\partial \theta_0} = -\frac{\frac{\partial F(\theta_1|\theta_0)}{\partial \theta_0}}{f(\theta_1|\theta_0)} \triangleq I_1(\theta_0, \theta_1).$$

The function $I_1(\cdot)$ is called the *impulse response function* in Pavan, Segal, and Toikka (2014) and it will play a key role in the following analysis. Since $\tilde{\theta}_0$ is independent of $\tilde{\varepsilon}_1$, the distribution of $\tilde{\varepsilon}_1$ does not vary as θ_0 changes. This fact implies that a characterization of the information rent of the buyer follows by the envelope theorem if $\partial F(\theta_1|\theta_0)/\partial \theta_0$ and $f(\theta_1|\theta_0)$ are sufficiently well-behaved. We refer to

$$(\tilde{\theta}_0, \tilde{\varepsilon}_1, F^{-1}(\varepsilon_1|\theta_0))$$

as the canonical (orthogonal) representation of the original information.

Since $\tilde{\theta}_0$ is independent of $\tilde{\varepsilon}_1$ by construction, we may view $\tilde{\theta}_0$ as the true private information to the agent and $\tilde{\varepsilon}_1$ as information not available at the moment of contracting. If $\tilde{\varepsilon}_1$

⁸If $F(\cdot|\theta_0)$ is constant for some interval or if it has upward jumps, we can define $F^{-1}(\varepsilon_1|\theta_0) = \inf\{\theta_1|F(\cdot|\theta_0) \geq \varepsilon_1\}$.

⁹Since

$$\varepsilon_1 \triangleq F(F^{-1}(\varepsilon_1|\theta_0)|\theta_0),$$

the second equality in the formula follows by total differentiation with respect to θ_0 .

were publicly observable and contractible, then we would indeed be dealing with the static problem in the previous subsection. If the solution of the mechanism design problem with observable and unobservable $\tilde{\varepsilon}_1$ coincide, then the principal does not have to give any information rent to the agent in excess of that contained in $\tilde{\theta}_0$. We will discuss below when this conclusion holds and when it does not.

As long as $\tilde{\theta}_1$ is first order stochastically increasing in $\tilde{\theta}_0$, we have $I_1 \geq 0$. In the appendix, we provide a construction of the t -period impulse responses for a general Markov process $\{\tilde{\theta}_t\}_{t=0}^T$ and we record here that

$$I_t(\theta^t, x^{t-1}) \triangleq - \prod_{k=1}^t \frac{\partial F(\theta_k | \theta_{k-1}, x_{k-1})}{f(\theta_k | \theta_{k-1}, x_{k-1})},$$

where we allow the transition function to depend on the current state and the current allocation, thus a controlled stochastic process. Since we will make use of an envelope theorem, we assume from now on that the impulse response functions are bounded.

A very rough protocol for solving dynamic mechanism design problems can now be given as follows. First, find the dynamic equivalent of the envelope formula (5) in the payoff 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 (6). Third, verify that the obtained solution satisfies the dynamic equivalent of the full incentive compatibility requirement (7). We pursue this program in the next few subsections.

4.3 Dynamic Payoff Equivalence

Since the dynamic mechanism design problem inherits the trade-off between

efficiency and information rent from the static problem, we must find a characterization for the information rent in terms of the allocation as in the static payoff equivalence theorem. This is where the above orthogonalized model becomes useful. In a model of perfect commitment, the mechanism designer maximizes her payoff from the perspective of $t = 0$. The orthogonalization gives a tractable solution for the agent's equilibrium payoff in $t = 0$ (and hence also her expected transfers).

Recall that in a dynamic direct mechanism (\mathbf{x}, \mathbf{p}) the agent reports her type θ_t at each t . Any allocation rule x under truthful reporting induces a stochastic process whose transitions are given by:

$$\theta_{t+1} \sim F(\cdot | \theta_t, x_t),$$

and we denote this process by $\lambda[\mathbf{x}]$. By the dynamic revelation principle in Sugaya and Wolitzky (2017), it is without loss of generality to consider a dynamic direct mechanism where the buyer reports her type θ_t truthfully on equilibrium path in each t , and any such mechanism is said to be incentive compatible.

We want to compute the equilibrium payoff to the agent with initial private information θ_0 :

$$V(\theta_0) = E^{\lambda[\mathbf{x}]} \left[\sum_{t=0}^T \delta^t (v(x_t(\theta_t, h_t), \theta_t) - p_t(\theta_t, h_t)) \right],$$

where the expectation is taken with respect to $\lambda[\mathbf{x}]$. The following theorem is a special case of the characterization of local incentive compatibility in theorem 1 of Pavan, Segal, and Toikka (2014).¹⁰

¹⁰Modulo changing the notation for a different starting date and starting history, the same characterization holds for $\partial V_s(\theta_s, \theta^{s-1}) / \partial \theta_s$ for all s and all $\theta^s \in \Theta^s$.

THEOREM 3 (Dynamic Revenue Equivalence): *If (\mathbf{x}, \mathbf{p}) is incentive compatible, then $V_0(\theta_0)$ is Lipschitz continuous and has almost everywhere the derivative:*

$$(9) \quad V'(\theta_0) = E^{\lambda^{|\mathbf{x}|}} \left[\sum_{t=0}^T I_t(\theta^t, x^{t-1}) \delta^t \frac{\partial v_t(x_t(\theta_t, h_t), \theta_t)}{\partial \theta_t} \right].$$

The reason for using the canonical representation is that it allows for an application of the Milgrom–Segal envelope theorem more easily than the original formulation. The representation also shows that the equilibrium payoff to the agent from truthful reporting is the same in the original game and the game where the $\tilde{\varepsilon}_t$ are publicly observed. This follows immediately from the assumption of truth telling and the envelope formula. We will return to this issue, but it should be noted already here that there are allocations where truth telling is not optimal in a mechanism with privately observed $\tilde{\varepsilon}_t$ but where the allocation rule can be implemented with observable $\tilde{\varepsilon}_t$.

As in the static case, the payoff equivalence theorem shows that the (dynamic) allocation pins down the agent’s payoff, and therefore her transfers, up to a constant. 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 (9) then becomes:

$$V'_0(\theta_0) = E \sum_{t=0}^T I_t(\theta^t, x^{t-1}) \frac{\partial U}{\partial \theta_t}.$$

The impulse response function measures the effect of a small change in θ_s on θ_t and $\partial U / \partial \theta_t$ measures the induced change in period t utility. All the other effects across

periods depend on the reported types, not the true types. The transfers that support the indirect utility of the agent can then be derived just as in the static model.

4.4 Dynamic Incentive Compatibility

Characterizing incentive compatible dynamic mechanisms is not hard at a very abstract level. A counterpart for the static integral monotonicity condition in equation (7) can be given as follows. A mechanism (\mathbf{x}, \mathbf{p}) is incentive compatible if the transfers are computed from formula (9) and if for each s, θ^s , and and report m_s in period s ,

$$(10) \quad \int_{m_s}^{\theta_s} \left[\frac{\partial V_s(q, \theta^{s-1})}{\partial \theta_s} - \frac{\partial V_s^{m_s}(q, \theta^{s-1})}{\partial \theta_s} \right] dq \geq 0,$$

where $\partial V_s^{m_s}(q, \theta^{s-1}) / \partial \theta_s$ is the derivative of the continuation payoff with respect to the true type q given the reports $\theta^s = (m_s, \theta^{s-1})$ and given that all future reports are truthful.

Even though a condition for full incentive compatibility can be expressed in a concise form, it does not yield an easy characterization of the feasible mechanisms for the principal. Pavan, Segal, and Toikka (2014) offer stronger sufficient conditions for implementability that are easier to verify. A very strong sufficient condition is that the payoff functions are supermodular in the allocation and type and the allocation (at all histories) is nondecreasing in all types.

It should also be noted that while the canonical representation is convenient for proving the payoff equivalence theorem, it is not helpful for analyzing full incentive compatibility. When one writes the payoff functions in terms of the canonical representation $(\theta_0, \{\tilde{\varepsilon}_t\}_{t=1}^T, \{Z_t\}_{t=1}^T)$ as defined in the appendix using the orthogonal shocks $\tilde{\varepsilon}_t$, one obtains:

$$\begin{aligned} v_t(x_t, \theta_t) &= v(x_t, Z_t(\theta_0, \varepsilon^t, x^{t-1})) \\ &\triangleq \hat{v}_t(x_t, \varepsilon^t, x^{t-1}). \end{aligned}$$

In other words, the agent's utility in period t depends on variables that are determined by past and current reported types and the whole sequence of realized ε_t . The first dependence is present in the original model as well, but the second type breaks the Markovian nature of the agents' problem. This adds new difficulties into checking full incentive compatibility, since checking for optimal behavior after non-truthful messages is no longer easy.¹¹

4.5 Optimal Dynamic Mechanism

We proceed to describe the solution of the optimal dynamic mechanism design. Similar to the static case, we start by considering a relaxed problem where the only constraint for the problem is that the transfers are calculated from the payoff equivalence theorem.

4.5.1 Relaxed Problem

We denote the dynamic payoff to the mechanism designer 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 designer's payoff as the difference between the social surplus and the agent's information rent. Using formula (9) to substitute for the payments gives, after the usual integration by parts, the following program for maximizing the designer's payoff from period $t = 0$ perspective:

¹¹Within the Markovian setting of the original model, the future incentives for truth telling are independent of past types. Hence if a report is optimal on the equilibrium path for type θ_t , it will also be optimal following non-truthful reports.

$$(11) \quad \max_{(\mathbf{x}, \mathbf{p})} E^{\lambda[\mathbf{x}]} \sum_{t=0}^T \delta^t \left(v(x(\theta_t, h_t), \theta_t) - c_t(x(\theta_t, h_t)) \right) - E^{\lambda[\mathbf{x}]} \frac{1 - F_0(\theta_0)}{f(\theta_0)} \times \left[\sum_{t=0}^T \delta^t I_t(\theta^t, x^{t-1}) \frac{\partial v(x(\theta_t, h_t), \theta_t)}{\partial \theta_t} \right] - V_0(\underline{\theta}_0).$$

The maximization is also subject to period 0 participation constraints:

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

We denote the first line in the objective function by $E^{\lambda[\mathbf{x}]}[S(x, \theta)]$ 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. If the stage payoff functions of the agent are supermodular, and if $\tilde{\theta}_{t+1}$ is first order stochastically increasing in $\tilde{\theta}_t$, then the individual participation constraint typically bind at the optimum for the lowest type and thus $V(\underline{\theta}_0) = 0$.

Even though the relaxed problem can be written rather concisely, solving

$$(12) \quad \max_{\mathbf{x}} E^{\lambda[\mathbf{x}]} \left[S(x, \theta) - \frac{1 - F(\theta_0)}{f(\theta_0)} \sum_{t=0}^T \delta^t I_t \frac{\partial v_t}{\partial \theta_t} \right]$$

involves dynamic programming and is not easy, in general. When discussing the applications below, we shall see some instances where more or less explicit solutions to the problem exist.

4.5.2 Properties of the Solution to the Relaxed Problem

If the process of $\{\tilde{\theta}_t\}_{t=1}^T$ does not depend on the allocations x^t and if there are no intertemporal restrictions on x_t , then a point-wise solution is possible as in the static case. Examples of this setting were covered already in Baron and Besanko (1984) and Besanko (1985). In this case, we can deduce some immediate properties of equation (12).

First, if the type is perfectly persistent, then $I_t(\theta^t) = 1$ for all t and θ^t . This implies that the optimal point-wise solution collapses to the static solution to

$$(13) \quad \max_x \left\{ s_t(x, \theta) - \frac{1 - F(\theta_0)}{f(\theta_0)} \frac{\partial v_t}{\partial \theta_t} \right\}.$$

Notice that here the distortions in the allocation rule remain over time.

Second, if the $\tilde{\theta}_t$ are independent of $\tilde{\theta}_0$ for $t > 0$, then x_t maximizes the social surplus for all $t > 0$. In this case, initial private information has no effect on future types.

Third, if the type process follows an AR(1) process:

$$\theta_t = \lambda \theta_{t-1} + \varepsilon_t,$$

with $\theta_{-1} = 0$, then one finds from the moving average representation,

$$\theta_t = \sum_{s=0}^t \lambda^{t-s} \varepsilon_s,$$

that $I_t(\theta^t) = \lambda^t$. Since we require that $|I_t(\theta^t)| < \infty$, we must have $\lambda \leq 1$. If we have a persistent random walk, i.e., $\lambda = 1$, then the solution is as with persistent types. If $\lambda < 1$, then the distortions from the efficient allocation vanish as $t \rightarrow \infty$. This simply reflects the fact that as time goes on, the effect of the initial shock $\theta_0 = \varepsilon_0$ on θ_t vanishes and at the moment of contracting, the principal and the agent have almost identical beliefs about θ_t for large t .

4.5.3 Full Solution

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 enough detail to allow us to verify sufficient conditions for full incentive compatibility.

The optimal solution in the original mechanism design problem coincides with the optimal solution in the modified problem of the canonical representation where the orthogonal shocks $\tilde{\varepsilon}_t$ are publicly observable. By the payoff equivalence theorem, this implies that the agent gets the same expected payoff in the two problems. Hence, it is reasonable to say that the agent does not benefit from the additional private information that she gets during the game.

At the same time, if the solution to the relaxed problem is not fully incentive compatible, this is no longer true. In the appendix, we present an example showing that the solution of the relaxed problem may not be incentive compatible even though the solution is incentive compatible when the $\tilde{\varepsilon}_t$ are publicly observable.

4.5.4 Implementing the Solution

Unfortunately there is no general recipe along the lines of the taxation principle for natural indirect implementations of the optimal direct mechanism in the dynamic setting. In some cases, the solution to the optimal contracting problem is suggestive of natural ways to implement the solution. For example, the sequential screening problem and the dynamic auction formats discussed below have solutions that can be implemented through option contracts and through various types of handicapped auctions.

5. Leading Applications

We shall now discuss how the general insights translate into specific solutions in a number of important economic applications.

5.1 Sequential Screening

Starting with Courty and Li (2000), the simplest model of bilateral trading with a dynamic flow of information has been called the *sequential screening model*. The canonical model extends over two periods, $t \in \{0, 1\}$, with trade taking place only in period $t = 1$. An uninformed seller proposes a mechanism to an informed buyer with type $\theta_0 \in [\underline{\theta}, \bar{\theta}]$ in $t = 0$. Her second period type $\theta_1 \in [\underline{\theta}, \bar{\theta}]$ is unknown to both parties at the moment of contracting, but it is common knowledge that its conditional distribution is given by $F(\theta_1 | \theta_0)$. The prior on θ_0 is denoted by $F_0(\theta_0)$.

The key economic question for the model is whether the seller can use the dynamic nature of information arrival to increase her expected revenue. Obviously she can sell using an optimal Myerson mechanism based on either $\tilde{\theta}_0$ or on $\tilde{\theta}_1$. The main observation in sequential screening models is that she can do strictly better by offering an option contract. In $t = 0$, the seller offers a menu of strike prices for period $t = 1$. Based on different θ_0 , the buyers choose different strike prices (obviously at different up-front payments) and this improves the seller's payoff.

Different interpretations are possible for θ_0 . It can be thought as a prior mean for θ_0 or, alternatively, it can be thought of as a measure of the precision of the agent's prior information about θ_1 . In the first case, it would be natural to assume that $\tilde{\theta}_1$ is first order stochastically increasing in $\tilde{\theta}_0$ while in the second, one would expect $\tilde{\theta}_1$ to be second order stochastically increasing in $\tilde{\theta}_0$.

Since trading takes place only in period $t = 1$, there is no loss of generality in

assuming that θ_1 is the value of the buyer in $t = 1$. We also assume that the good is indivisible (or alternatively, the payoffs are linear in quantities) and that the seller has no value for the object herself. This leads to the payoffs:

$$u_S(\theta_1, x, p) = p,$$

$$u_B(\theta_1, x, p) = \theta_1 x - p,$$

for the seller and the buyer respectively, where x is the probability of trading and p is the transfer from the buyer to the seller. We also assume that the outside option for the buyer yields payoff 0.

We can use the general result from the previous section to see rather quickly how to arrive at this solution. A direct dynamic mechanism is now a pair of functions:

$$x : \Theta_0 \times \Theta_1 \rightarrow [0, 1],$$

$$p : \Theta_0 \times \Theta_1 \rightarrow \mathbb{R}_+.$$

We recall that the single period impulse response function can be written as:

$$I_1(\theta_0, \theta_1) = -\frac{\frac{\partial F(\theta_1 | \theta_0)}{\partial \theta_0}}{f(\theta_1 | \theta_0)}.$$

Together with equation (9), this gives:

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

Solving for the expected transfer (i.e., the seller's expected payoff), the relaxed problem becomes (after an integration by parts and using the individual rationality constraint $V(\underline{\theta}) = 0$):

$$\max_x \int_{\Theta_0} \int_{\Theta_1} x \left[\theta_1 - \frac{1 - F_0(\theta_0)}{f_0(\theta_0)} I_1(\theta_0, \theta_1) \right] \times f(\theta_1 | \theta_0) f(\theta_0) d\theta_1 d\theta_0.$$

But this is a linear problem in x , hence the relaxed solution is easy to find.

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

$$\psi(\theta_0, \theta_1) \triangleq \theta_1 - \frac{1 - F_0(\theta_0)}{f_0(\theta_0)} I_1(\theta_0, \theta_1).$$

This modifies the classic Myersonian virtual value by multiplying the information rent component $(1 - F_0(\theta_0))/f_0(\theta_0)$ by the impulse response I_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 4 (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*

$$x(\theta_0, \theta_1) = \mathbf{1}_{\{\theta_1 \geq q(\theta_0)\}},$$

and the transfer is computed from the envelope formula.

As anticipated, in the optimal mechanism the buyer pays an up-front fee $p(\theta_0)$ for 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.

Courty and Li (2000) show that in the case where $\tilde{\theta}_1$ is second order stochastically increasing in $\tilde{\theta}_0$, the standard Myersonian downward distortions may be reversed. If $\tilde{\theta}_1$ is first order stochastically increasing in $\tilde{\theta}_0$, this is not possible. This result can be understood in terms of the sign of the impulse response function in the two cases. Under first-order stochastic dominance (FOSD), $I_1(\theta_0, \theta_1)$ is always positive. For the case of second-order stochastic dominance (SOSD), it may well be negative, thus leading to a reversal in the direction of the distortions.

Eső and Szentes (2007) extend the model to allow for multiple bidders for the good (otherwise the model is identical to the model above). They find an optimal auction—called the *handicap auction*—where the bidders can make up-front payments in the first period to influence the allocation rule determining the second period allocation (the handicaps for the final auction). In order to analyze the model, Eső and Szentes (2007) introduced the orthogonalization process described in section 4.2. They compare the revenue to the seller under two scenarios: one where she releases the orthogonal signals to the buyers and one where she does not. They conclude that the seller is always better off releasing the information.

Bergemann and Wambach (2015) and Li and Shi (2017) offer extensions of the sequential screening model that incorporate information *and* mechanism design. Li and Shi (2017) show that even though the seller always wants to release all of the orthogonalized information to the buyer, she may prefer to send garbled information based on the original (not orthogonalized) type θ_1 . The question of what types of disclosure policies are optimal in this setting is still open.

5.2 Selling Options

The second illustrative example is a stopping problem, rather than a selling problem at a fixed deadline. 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 must decide to whom to allocate the object. Suppose the evolution of the willingness to pay by bidder i is given by:

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

with $\theta_{i,0} \sim G_i(\theta_{i,0})$, $\varepsilon_{i,t} \sim H_i(\cdot)$, 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,

$$V_{i,0}(\theta_{i,0}) = E \sum_{t=0}^T \delta^t \frac{1 - G_i(\theta_{i,0})}{g_i(\theta_{i,0})} \gamma^t x_{i,t}(\theta),$$

and find that the expected revenue to the seller is

$$E \sum_{t=0}^T \sum_{i=1}^N \delta^t \left[\theta_{i,t} - \frac{1 - G_i(\theta_{i,0})}{g_i(\theta_{i,0})} \gamma^t \right] x_{i,t}(\theta).$$

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

$$\max_i \left\{ \theta_{i,t} - \frac{1 - G_i(\theta_{i,0})}{g_i(\theta_{i,0})} \gamma^t \right\}$$

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.

5.3 Bandit Auctions

A single indivisible object is allocated in each period among n possible bidders who

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 $x_{i,t} = 0$, then $\theta_{i,t+1} = \theta_{i,t}$, if $x_{i,t} = 1$, then

$$(14) \quad \theta_{i,t+1} = \theta_{i,t} + \varepsilon_i(n_i(t)),$$

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 version of the process outlined in section 2.2, the number of observations from the process (here $n_i(t)$) and the current posterior mean (here $\theta_{i,t}$) form a sufficient statistic. We can interpret the allocation process as intertemporal licensing where the current use of the object is determined by 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:

$$(15) \quad \prod_{t=r}^s \left(- \frac{\partial F_i(\theta_{i,t+1} | \theta_{i,t})}{\partial \theta_{i,t}} \frac{1}{f_i(\theta_{i,t+1} | \theta_{i,t})} \right) = 1.$$

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

$$\max_{x \in X} E \sum_{t=0}^T \sum_{i=1}^N \delta^t \left[\theta_{i,t} - \frac{1 - F_i(\theta_{i,0})}{F_i(\theta_{i,0})} \right] x_i(\theta_{i,t}),$$

where $X = \{(x_1, \dots, x_N) \in \mathbb{R}_+^N \mid \sum_i x_i = 1\}$. 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, which is determined exclusively by the initial private information $\theta_{i,0}$. Moreover, by (15), the impulse response function, and hence the handicap, is constant in time and determined only by 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.

5.4 Repeated Sales

A common economic setting where long-term contracts govern the interaction between buyer and seller is the repeated sales problem. The buyer anticipates that he might purchase a good or a service repeatedly over time, but is uncertain about his future valuation of the good. At any point in time, his willingness to pay is private information, and the current willingness to pay is a good prediction of the future willingness to pay. A variety of dynamic contracts are used to support the provision of services, as documented by DellaVigna and Malmendier (2006), Grubb (2009), and Eliaz and Spiegler (2008) for gym memberships, mobile phone contracts, and many other services.

These allocation problems can be viewed as being separable across periods in two important aspects: (i) the set of feasible

allocations at time t is independent of the history of the allocations, and (ii) the flow utility function depends only on current type. This class of models is particularly tractable since a point-wise solution to the relaxed problem is quite easily obtained and the conditions for full incentive compatibility can be directly checked. In fact, the earliest contributions to the dynamic mechanism design literature, Baron and Besanko (1984) and Besanko (1985) restricted attention to time-separable problems of this form.

Bergemann and Strack (2015) consider time-separable allocation problems in continuous time. They leverage the structure of the continuous-time setting to obtain closed-form solutions of the optimal contract. In the leading example of repeat sales of a good or service, they establish that many commonly observed contract features such as flat rates, free consumption units, and two-part tariffs can emerge naturally as part of the optimal contract.

In their setting, the flow value is given by $v_t x_t - p_t$, v_t is the willingness to pay in time t , and x_t the quantity or quality assigned to buyer in period t . The willingness to pay is assumed to be a function,

$$v_t = \phi(t, \theta_0, W_t),$$

that is weakly increasing in the initial type θ_0 and the value of a Brownian motion W_t in period t . With the time separability of the allocation across periods, the virtual utility in period t is simply given by

$$(16) \quad v_t - \frac{1 - F(\theta_0)}{f(\theta_0)} \frac{\partial \phi(t, \theta_0, W_t)}{\partial \theta_0}.$$

This is simply the continuous-time analogue to the relaxed problem that we derived earlier in (13), where the derivative

$$\frac{\partial \phi(t, \theta_0, W_t)}{\partial \theta_0},$$

often referred to as the *stochastic flow*, replaces the product of the marginal flow value times the impulse response function. The nature of the initial information θ_0 together with the shape of the stochastic process now determine how the stochastic flow, and ultimately the optimal allocation, vary over time.

Bergemann and Strack (2015) then analyze how the optimal contract depends on the nature of initial private information and the structure of the stochastic process. In their leading case, the valuation evolves as a geometric Brownian motion

$$dv_t = (v_t - \underline{v})\sigma dW_t,$$

where $\underline{v} \geq 0$ is a lower bound on the flow utility, and W_t is a Brownian motion. If the initial private information is simply the initial value of the process, $\theta_0 = v_0$, then the stochastic flow is simply

$$\frac{\partial \phi(t, \theta_0, W_t)}{\partial \theta_0} = \frac{v_t - \underline{v}}{v_0}.$$

Thus, the corresponding expression from discrete time, the impulse response function, reduces to a simple expression. They can consequently show that a menu of flat-rate contracts, different two-part tariffs, or different free minute contracts can arise as optimal solutions, depending on the value of the lower bound \underline{v} and the flow cost of providing the service, given by $c(x)$.

By contrast, if the initial private information θ_0 is the drift of the geometric Brownian motion, thus

$$dv_t = v_t(\theta_0 dt + \sigma dW_t),$$

then the stochastic flow can be computed to be

$$\frac{\partial \phi(t, \theta_0, W_t)}{\partial \theta_0} = v_t t.$$

Now the optimal contract is a menu of leasing contracts with deterministic deadlines as the flow virtual utility takes the form:

$$v_t \left(1 - \frac{1 - F(\theta_0)_t}{f(\theta_0)} \right).$$

Interestingly, the distortion is linearly increasing in time. It follows that in contrast to many of the models analyzed so far, the allocative distortion is now increasing over time rather than decreasing over time. A noteworthy aspect of this last example is that the initial private information of the agent is not the initial value of the stochastic process, but rather a parameter of the stochastic process itself.

5.5 Private Information about the Stochastic Process

In fact, a number of recent contributions have 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

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.

5.6 Beyond the First-Order Approach

The method of analysis for the dynamic contracting problem above relies heavily on the payoff equivalence theorem, also known as the first-order approach. For this approach to be successful, the models must be such that the solutions to the relaxed problem are incentive compatible. Battaglini and Lamba (2017) show that in models with discrete types, the first-order approach becomes problematic if the agents interact frequently. In particular, the solution to the first-order problem is no longer monotonic if types are highly persistent. They propose and analyze optimal contracts in the class of strongly monotonic allocation functions and show that these contracts are approximately efficient in the class of all incentive compatible contracts.

Garrett, Pavan, and Toikka (2018) take a different approach to the problem. They characterize necessary properties of the optimal contract by a relatively simple perturbation argument. They show that regardless of whether the first order approach is applicable or not, the optimal contract must have vanishing distortions as long as the underlying process on types is sufficiently mixing, in the sense that the impact of initial information on future types vanishes. Hence this paper confirms, for a larger class of models, one of the key findings in Battaglini (2005) derived for models with binary types.

6. Dynamic Populations

In this section, we consider mechanism design problems where the population of privately informed agents changes over time. To fix ideas, we return to our leading example of a seller who has a fixed capacity K of indivisible goods to sell by a (possibly infinite) deadline T . Potential buyers arrive according to a stochastic process and the

seller wants to extract as much revenue as possible from them. Variants of this problem have been studied in the literature on *revenue management* in management science and in operations research.

Important economic examples fit this description very nicely. By far the most important and most analyzed example is the pricing of airline tickets. As airlines customers have noticed, prices for identical tickets on a given flight vary over time. The airlines industry uses various dynamic pricing and allocation methods for the seats. They use time-varying posted prices that may depend on the query data for the flight in question, and sometimes also more complicated mechanisms allowing for the possibility of securing a future price by paying an up-front fee. These features are important for potential buyers as well. Forward-looking buyers should time their purchases optimally given their expectation of the price path. In fact, services such as Kayak have been developed to alert buyers to particularly good moments to purchase tickets.

Optimal timing of the purchases is a natural element in any dynamic mechanism design problem of this type. The findings in this literature have a clear substantive message. In a wide class of models, the literature shows that sellers cannot be made worse off if the buyers are forward looking rather than myopic. Moreover, the analysis of models with forward-looking buyers guides practical implementations for the revenue maximizing scheme.

A key modeling decision with dynamic populations is whether the arrival of a buyer is publicly observable or not. We start with the simpler models where observability is assumed. We discuss also models where arrival is private information to the buyer. In this case, the buyer's type has two dimensions: age and valuation, but the model has enough structure that the analysis remains tractable.

6.1 *Observable Arrival of Short-Lived Buyers*

The first approach in the revenue management literature was to assume away problems of asymmetric information, i.e., assume that the seller observes the valuations and that the buyers are short lived (i.e., they disappear if they are not allocated the good).¹³ In this framework, the seller's problem coincides with the problem of surplus maximization and the analysis is a standard (but not analytically simple) exercise in dynamic programming. Within the revenue management literature, the key analytical aspect of the problem is to find a characterization for the optimal allocation rule for the goods as a function of remaining objects k at any point in time and remaining time t to the deadline. The key finding in this literature is that the optimal allocation rule is often given by a cutoff rule in the set of types or valuations: allocate at (k, t) if and only if $v_t \geq g(k, t)$ for some function g that is typically decreasing in k and increasing in t .¹⁴

The mechanism design approach to this problem emphasizes the effects of incentive compatibility when the buyers's types are not observable to the seller. When buyers are short lived and their process of arrivals is observable to the seller, then we are specializing the general model to the case described at the end of subsection 2.2, where the agent's type in period t is $\theta_{i, \tau_i} \in [0, \bar{\theta}]$ if $t = \tau_i$ for some (possibly random) publicly observed arrival period τ_i and $\theta_{i, \tau} \in \theta_0$ if $t \neq \tau_i$ and her payoff from allocation x in a (possibly random) period τ_i is $v_i(x_{i, \tau_i}, \theta_{i, \tau_i})$.

¹³The classical references are Derman, Lieberman, and Ross (1972) for the case of known distribution of buyer valuations and Albright (1977) for the case where the seller learns about the distribution based on the observed types.

¹⁴With an infinite deadline, the problem becomes stationary if the arrival process of buyers is stationary and the solution of the process simplifies considerably.

In the simplest case, the buyers have unit demands for the object and they have independent valuations. By the payoff equivalence theorem, incomplete information about the type of the buyer transforms the maximization of total expected revenue to the relaxed problem of maximizing expected virtual surplus. For notational convenience, we assume that the distribution of the realized type does not depend on the arrival time τ_i , but this could be easily accommodated in the model as well.

For the case of identical objects and with $v_i(x_{i, \tau_i}, \theta_{i, \tau_i}) = \theta_{i, \tau_i} x_{i, \tau_i}$, where $x_{i, \tau_i} \in \{0, 1\}$ indicates whether the object is allocated or not to i in τ_i , we can write the expected revenue of the seller in terms of the expected virtual utility:

$$E_{\theta} \sum_{i=1}^N \delta^{\tau_i} \left[\theta_i - \frac{1 - F(\theta_i)}{F(\theta_i)} \right] x_{i, \tau_i}$$

such that

$$\sum_{i=1}^N x_{i, \tau}(\theta_i) \leq K \quad \text{for all } \theta_i \in \Theta_i,$$

where the expectation is taken over the vector θ of type processes and the allocation decisions depend only on the realized part θ^{τ_i} of the process at τ_i .

Full incentive compatibility typically boils down to an appropriate monotonicity requirement for the allocation rule in type $\theta_{i, t}$. For the case of identical objects, monotonicity is equivalent to a cutoff characterization of the allocation rule. Hence, a sequence of posted prices can always implement the optimal allocation. Versions of this problem have been analyzed in a sequence of papers by Gershkov and Moldovanu (2009a, b), and additional results collected in Gershkov and Moldovanu (2014). The revenue maximization problem in Gershkov and Moldovanu (2009a) allows for the possibility that the K objects to be allocated have

different (vertical) qualities. Gershkov and Moldovanu (2009b) shows that learning or correlation in types may cause problems for the monotonicity even for the socially efficient allocation rule. Thus, the difficulty of obtaining a monotone allocation rule does not arise solely due to non-monotonic virtual surpluses. This observation can be illustrated nicely within our leading example.

Example 3 (Learning and Incentive Compatibility)

A single indivisible good is to be allocated efficiently to one of two bidders $i \in \{1, 2\}$, both with a strictly positive valuation for the object. It is commonly known that $\tau_i = i$ and therefore the relevant social allocation decisions are given by $x \in \{1, 2\}$, where the first choice indicates allocating the object to bidder 1 in period 1 and the second indicates allocating to bidder 2 in period 2. The planner's objective is to maximize the social surplus and she has a discount factor δ .

The valuation $\theta_1 \in [0, 1]$ of bidder 1 is known at the outset of the game. Bidder 2 learns her value in period 2. The valuations can come from one of two possible distributions: θ_i is uniformly distributed on either $[0, 1/2]$ or $(1/2, 1]$. The prior probability of each of these distributions is identical.

With observable types, the planner's optimal solution is immediate: allocate to agent 1 in period 1 if and only if

$$\theta_1 \in \left[\frac{1}{2}\delta, \frac{1}{2} \right] \cup \left[\frac{3}{4}\delta, 1 \right].$$

As long as $\delta > 2/3$, this allocation rule is not monotone, and since the bidders' payoffs are supermodular, it fails to be incentive compatible for the case of unobserved types. If a bidder can make a payment only in the period when he receives the good, we see that there is no way of implementing the efficient decision rule in the model with incomplete information. This problem does not arise if we can condition payments on the reports

of both types. The team mechanism derived in section 3 works nicely here if such contingent payments are allowed. Hence the example points out the problems that arise as a consequence of the (often quite realistic) requirement that monetary transfers occur only in conjunction with physical allocation decisions. This requirement sometimes goes under the name of "online" payments.

6.2 Unobservable Arrival of Long-Lived Buyers

With unobservable arrivals, the buyers will have an incentive to time their purchases strategically. If prices decrease over time, they will delay reporting to the mechanism in order to get a better deal later on. Any incentive compatible mechanism must take this possibility into account. By contrast, if the arrivals are publicly observed then this is not a concern. The seller may simply commit not to offer any contracting opportunities except in the period of arrival. If the arrivals are not observed, the seller cannot distinguish between new arrivals in any period t from those that arrived earlier and waited with their announcement.

Board and Skrzypacz (2016) consider the sales problem of K identical indivisible units to a population of arriving buyers when the arrivals are private information. In their model, the statistical properties of the arrivals and valuations are common knowledge at the beginning of the game and arrivals and types satisfy independence across buyers and across periods. They show that the optimal selling mechanism is surprisingly simple: it is a deterministic sequence of posted prices depending on (k, t) . Interestingly, this pattern leads to waiting by the buyers along the equilibrium path. Even if the seller knew the past realized arrivals, this would not change the solution. If the demand is decreasing over time, they find a surprisingly explicit analytical solution for the problem. The other main substantive finding is that

under the optimal mechanism, the seller gets a higher revenue than she would get if the buyers were short lived. This happens even though sales are more back loaded in the case with forward-looking buyers and prices falling. If the modeling assumptions are relaxed, each of these two main predictions may fail and complete solutions seem difficult to obtain. See, for example, the analysis in Mierendorff (2016) for the case where the agents are forward looking but may disappear or have different discount factors.

Gershkov, Moldovanu, and Strack (2018) extend the model to cover the case where the buyers' arrival process is initially not known. This relaxes the assumption of independence in arrival times. They show that even though the optimal allocation is no longer implementable through anonymous posted prices, a simple name-your-own-price mechanism can be used as an indirect mechanism that achieves the maximal revenue. Further results in the paper show that the seller does not benefit from hiding information, say, about the existing stock of k units and that forward-looking buyers still benefit the seller as in Board and Skrzypacz (2016).

Since the models of dynamic populations where the buyers' types, the willingness to pay, are fixed over time lead to considerations of strategic timing, it is natural to ask how the case where buyers' types change over time would change the problem. In this case, there are two reasons for optimizing over the purchasing time: the price may be more favorable in the future or the type may change to one with a higher information rent.

A related issue 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.

Garrett (2017) considers a model of sales of a nondurable good where the buyer appears at a random future time, the arrival is private information, and the buyer's privately known valuation changes over time. He shows that in the optimal mechanism, the principal commits to punishing the agent for late arrival by inducing more inefficiencies to diminish information rents from manipulating the entry time.

A more ambitious attempt in this direction appears in Garrett (2016), where generations of new buyers are arriving over time to contract with a seller. A full mechanism design approach is not tractable in this case, and the paper restricts attention to an optimal time-dependent sequence of posted prices. Using anonymous posted prices implies that old and new buyers with the same valuation type have the same incentives for all purchases. In this sense, explicit penalties for late arrivals are not possible. In an otherwise stationary environment, the optimal posted price fluctuates. This comes as a surprise after the well-known result in Conlisk, Gerstner, and Sobel (1984) showing that stationary prices are optimal because the forward-looking buyers with high values anticipate lower prices and are therefore reluctant to buy at high prices. In Garrett (2016), high valuation buyers are more keen to buy immediately, since they understand that their type may decrease in the future.

Bergemann and Strack (2019) analyze a dynamic revenue-maximizing problem in continuous time when the arrival time of the agent is uncertain and unobservable to the seller. The valuation of the agent is private information, as well, and changes over time. They derive the optimal dynamic mechanism, characterize its

qualitative structure, and derive a closed form solution. As the arrival time of the agent is private information, the optimal dynamic mechanism has to be stationary to guarantee truth-telling. The truth-telling constraint regarding the arrival time can be represented as an optimal stopping problem. 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.

7. Connections to Nearby Models

In this section, we briefly discuss 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. 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. Whether private information exists at the moment of contracting or not is not that important for this literature, since the incentives–insurance trade-off emerges in any case as private information is generated.

In models of financial economics, a key assumption is that the privately informed managers may be risk-neutral but that they do not have sufficient funds to buy the entire enterprise. This is typically formalized through a limited liability constraint stating that the manager (the agent) cannot make payoffs to the owner (the principal). Recent work starting with Clementi and

Hopenhayn (2006), DeMarzo and Sannikov (2006), DeMarzo and Fishman (2007), and Biais et al. (2007) has analyzed the problem of incentivizing a manager who privately observes the cash flow of a firm.

7.1 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 u_0 (e.g., to satisfy a participation constraint). 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. This problem has attracted a large amount of attention starting with Green (1987) and Thomas and Worrall (1990).¹⁵ Our goal here is not to assess this literature, but merely point out how it connects to the models in the previous sections of this survey.

The consumer derives utility $v(x_t, \theta_t)$ in period t from allocation x_t if her type is θ_t . We consider incentive compatible dynamic direct mechanisms. The agent's problem is typically formulated as a dynamic programming problem (induced by the mechanism):

$$V(\theta_t) = \max_t \left\{ v(x(r_t, h_t), \theta_t) + E[V(\tilde{\theta}_{t+1}) | \theta_t] \right\}.$$

Under sufficient regularity conditions, value functions V_t satisfying these equations exist and are sufficiently well behaved for an application of the envelope theorem. Indeed, if one assumes that the set of possible types is a connected interval and that the process of types has full support, then

¹⁵Many of the issues also arise in dynamic incentive provision models with hidden actions starting with Rogerson (1985).

an application of the envelope theorem yields:

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

Integration by parts gives:

$$V'(\theta_t) = \frac{\partial v(x_t, \theta_t)}{\partial \theta_t} + E \left[\frac{-\frac{\partial F(\theta_{t+1} | \theta_t)}{\partial \theta_t}}{f(\theta_{t+1} | \theta_t)} V'(\theta_{t+1}) \right].$$

Hence by iterating this formula forwards, one gets as before:

$$V'(\theta_0) = E \left[\sum_{t=0}^{\infty} I_t \delta^t \frac{\partial v(x_t, \theta_t)}{\partial \theta_t} \right].$$

In other words, a similar envelope theorem characterization for the agent's utility is still possible in this model. Many papers in the new public finance literature adopt this approach. For example Farhi and Werning (2007) study dynamic insurance schemes from this perspective.

But following with this first-order approach to dynamic problems, the next step of substituting the agent's payoff into the principal's objective unfortunately fails because the utility is not quasi-linear. 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.

7.2 Managerial Contracts and Hidden Actions

Garrett and Pavan (2012) consider a model where a risk-neutral principal contracts with

a risk-neutral manager whose type (productivity) changes over time. The manager (the agent) has to be incentivized to take the optimal action at each point in time and the distortions now refer to the dynamic distortions relative to the model where incentives are provided in a setting with no private information. The paper shows that as long as the impulse response functions in the model are positive (for the privately observed productivity of the manager), then the distortions to the incentives diminish over time and incentives become more high powered.¹⁶

Limited liability protection on the 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 the principal from selling the enterprise to the agent at the outset even when there is no initial asymmetric information and hence there are no losses due to information rent left with the agent. A canonical model for this literature is one where the agent reports a privately observed i.i.d. cash flow to the principal in each period. The mechanism determines the transfers to the agent and a probability of continuing the project as functions of the (history of) reported cash flows. A key finding in this literature is that over time, the contract becomes more efficient, i.e., the probability of inefficient liquidation decreases over time. It should be noted that the intuition for this finding is very different compared to the models surveyed above. With limited liability, the optimal contract effectively saves funds for the agent so that she can buy the enterprise at a later time. Since recent surveys of this large literature

¹⁶In contrast, Garrett and Pavan (2015) consider the case where a risk-neutral principal provides incentives for a risk-averse manager and shows that the power of incentives vanishes over time to reduce the overall riskiness in the contract.

exist (see for example Biais, Mariotti, and Rochet 2013), we do not survey the topic here.

8. *Concluding Remarks*

It was our objective to give a broad and synthetic introduction to 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. Still, many interesting questions remain wide open. We shall describe some of them in these final remarks.

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 into 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 or short-term arrangements rather than long-term contracts. An important question then is why more transactions are not 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 et al. (2018) 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 et al. (2016). But in contrast to the Bayesian approach most commonly taken by economic theorists who explicitly identify and design the optimal mechanism, theoretical computer scientists often describe achievable performance guarantees. The bounds are frequently achieved by mechanisms that have computational advantages in terms of computational complexity relative to the possibly unknown exact optimal mechanism.

As an important friction to long-term arrangements is presumably the diversity in expectations about future events among the players, it is natural to ask to what extent the relevant 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—somewhat

surprisingly for a dynamic perspective—have received much less attention. A dynamic mechanism requires voluntary participation at the ex ante and interim stages via 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 in the initial period. Interim participation constraints can be handled by allowing the agents to post bonds at the initial stage. In general, the mechanism does not provide any guarantees about ex post participation constraints.¹⁷ In fact, 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. More generally, 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).

An interesting set of issues arise when the mechanism itself can govern only some of the relevant economic transaction. A specific setting where this occurs is markets with resale. Here, the design of the optimal mechanism in the initial stage of the game is affected by the interaction in the resale market, see for example Calzolari and Pavan (2006); Dworzak (2017); Carroll and Segal (2016); and Bergemann, Brooks, and Morris (2017). In particular, the information that is generated by the mechanism may affect the nature of the subsequent interaction, and

¹⁷Many commonly observed dynamic contracts do in fact violate ex post participation constraints. For example, an insurance company does not return the premium in case of no accident.

thus the tools from information design and mechanism design may jointly yield interesting new insights.

APPENDIX A. DYNAMIC PIVOT MECHANISM AND INDEPENDENCE

To see why the restriction to independent values is necessary, recall the transfer rule for agent i in the static pivot mechanism:

$$p_i(\theta) = -\sum_{j \neq i} u_j(x^*(\theta), \theta_j) + \sum_{j \neq i} u_j(x^*_{-i}(\theta_{-i}), \theta_j),$$

where $x^*_{-i}(\theta_{-i})$ is the optimal allocation for agents different from i . In the static case, x^*_{-i} depends only on the vector θ_{-i} by the assumption of private values regardless of any statistical dependencies between the agents' types. In the dynamic case, with correlated values, $\theta_{i,t}$ might have an effect on $x^*_{-i,t}$ even when fixing $\theta_{-i,t}$. As a result, both sums on the right hand side of (2) depend on $\theta_{i,t}$ and this distorts the incentives for truthful reporting. The following example illustrates this point.

Example 4 (Capacity Allocation and Correlated Types)

Three agents, $i \in \{1, 2, 3\}$, are bidding for a single indivisible object over three periods. Let $x_t \in \{1, 2, 3\}$ denote the possible allocations to i in period t . If the good is allocated in period s , then $\theta_{i,t} = 0$ for all i and all $t > s$ (say because it is not worthwhile to pay a cost to learn the valuation for an object that was already sold). Assume also that $\theta_{i,t} = 0$ if $i \neq t$ to indicate that agent i is active at most in period $t = i$. The payoff to agent i in period t is then $\Pr\{x_t = i\} \theta_{i,t}$.

Assume that $\theta_{1,1} \in \{3, 3 - \varepsilon\}$, $\theta_{2,2} = 1$ if $x_0 = N$ and $\theta_{3,3} = 0$ otherwise, $\theta_{3,3} \in \{2, \varepsilon\}$ if $x_1 = x_2 = N$ and $\theta_{1,1} = 0$

otherwise. Let π denote the prior probability that $\theta_{2,2} = 2$ given that $x_0 = x_1 = N$, and assume that $2\pi\delta > 1$. Then it is optimal to have $x_1 = N$, $x_2 = 2$ conditional on $x_0 = N$. As long as $\varepsilon < 1$, the efficient decision is to have $x_0 = 0$. But it is clear that this is not incentive compatible if the transfers are calculated using the pivotal rule. To minimize the payments, agent 0 should always report type $\theta_{0,0} = 3 - \varepsilon$.

APPENDIX B. THE CANONICAL REPRESENTATION OF A MARKOV PROCESS

Consider an arbitrary Markov process $\{\tilde{\theta}_t\}_{t=0}^T$ with the prior distribution $F_0(\theta_0)$ and the transition kernel $\tilde{\theta}_{t+1} \sim F(\cdot|\theta_t)$ on Θ . Since $\tilde{\varepsilon}_{t+1} = F(\tilde{\theta}_{t+1}|\theta_t)$ is uniformly distributed for all θ_t , we can deduce θ_{t+1} from $(\theta_0, \varepsilon_1, \dots, \varepsilon_t)$ by using the recursive formula:

$$(17) \quad \theta_t = F^{-1}(\varepsilon_t|\theta_{t-1}) \triangleq Z_t(\theta_0, \varepsilon^t).$$

Since the functions $Z_t(\cdot, \cdot)$ are common knowledge at the start of the game, $(\tilde{\theta}_0, \dots, \tilde{\theta}_t)$ contains the same information as $(\theta_0, \tilde{\varepsilon}_1, \dots, \tilde{\varepsilon}_t)$. Pavan, Segal, and Toikka (2014) call the collection $(\tilde{\theta}_0, \{\tilde{\varepsilon}_t\}_{t=1}^T, \{Z_t\}_{t=1}^T)$ the *canonical representation* of the process $\{\tilde{\theta}_t\}_{t=0}^T$.

In computing the dynamic payoff equivalence formula, we need to evaluate the impact of the initial private information on future types. Using the chain rule, we can compute from equation (17) the impact $\hat{I}_t(\varepsilon^t)$ of the initial private information θ_0 on θ_t for any fixed sequence of orthogonalized future types ε^t :

$$(18) \quad \hat{I}_t(\varepsilon^t) = \prod_{k=1}^t \frac{\partial F^{-1}(\varepsilon_k|\theta_{k-1})}{\partial \theta_{k-1}} \\ = - \prod_{k=1}^t \frac{\frac{\partial F(\theta_k|\theta_{k-1})}{\partial \theta_{k-1}}}{f(\theta_k|\theta_{k-1})} \triangleq I_t(\theta^t).$$

The function $I_t(\cdot)$ that expresses this impact in terms of the original type formulation is called the *impulse response function* in Pavan, Segal, and Toikka (2014) and it plays a key role the characterization of the agent's information rent.¹⁸ Notice that as long as $\tilde{\theta}_t$ is first order stochastically increasing in $\tilde{\theta}_{t-1}$, we can show that $I_t(\theta^t) \geq 0$. From now on, we assume that $|I_t(\theta^t)| < K$ for some $K < \infty$ so that the formula for determining $I_t(\theta^t)$ makes sense, and we let $I_0(\theta_0) = 1$ for all θ_0 .

The canonical representation is by no means a unique representation of the original model in terms of initial information and subsequent independent information. The next appendix gives an example where a noncanonical representation allows us to overcome a differentiability problem in the canonical representation.

We have presented the construction here for homogeneous Markov processes, but the same procedure can be used to obtain a canonical representation for the more general processes $F(\theta_{t+1}|\theta_t, x^t)$. In this more general case, we denote the impulse response functions by $I_t(\theta^t, x^{t-1})$.

APPENDIX C. IMPLEMENTABILITY IN THE ORTHOGONALIZED MODEL

This example, communicated to us by Juuso Toikka, gives an example of a problem where the solution of the relaxed problem is not (fully) incentive compatible, but where it can be implemented if the orthogonalized information is publicly observable.

¹⁸Since

$$\varepsilon_k \equiv F(F^{-1}(\varepsilon_k|\theta_{k-1})|\theta_{k-1}),$$

the second equality in the formula follows by total differentiation with respect to θ_{k-1} :

$$\frac{\partial F^{-1}(\varepsilon_k|\theta_{k-1})}{\partial \theta_{k-1}} = - \frac{\frac{\partial F(\theta_k|\theta_{k-1})}{\partial \theta_{k-1}}}{f(\theta_k|\theta_{k-1})}.$$

A seller (mechanism designer) with cost function $c_i(x_1) = x_1^2/2$ sells to a privately informed buyer in period $t = 1$. In period $t = 0$, there is no trade, hence the allocation decision for that period is trivial. The buyer's (the agent's) type in $t = 0$ is uniformly distributed, $\theta_0 \sim U[0, 1]$, and in period $t = 1$, it remains unchanged with probability q . With probability $(1 - q)$, the type is drawn independently of θ_0 . The canonical representation of the model is $(\tilde{\theta}_0, \tilde{\varepsilon}_1, Z_1(\theta_0, \varepsilon_1))$, where

$$Z_1(\theta_0, \varepsilon_1) = \begin{cases} \frac{\varepsilon_1}{1-q} & \text{if } 0 \leq \varepsilon_1 < (1-q)\theta_0 \\ \theta_0 & \text{if } (1-q)\theta_0 \leq \varepsilon_1 \leq (1-q)\theta_0 + q \\ \frac{\varepsilon_1 - q}{1-q} & \text{if } (1-q)\theta_0 + q < \varepsilon_1 \leq 1 \end{cases}$$

Notice that $Z_1(\theta_0, \varepsilon_1)$ is not differentiable in θ_0 at $\theta_0 \in \left\{ \frac{\varepsilon_1}{1-q}, \frac{\varepsilon_1 - q}{1-q} \right\}$ and hence the envelope theorem is not directly applicable.¹⁹ This problem can, however, be overcome by selecting a different (noncanonical) orthogonal representation with a two dimensional $\eta_1 = (\eta_{11}, \eta_{12})$ where η_{11} is a Bernoulli random variable with $\Pr\{\eta_{11} = 1\} = 1 - \Pr\{\eta_{11} = 0\} = q$ and η_{12} is an independent uniform random variable. Then we can write

$$\theta_1 = \hat{Z}_1(\theta_0, \eta_{11}, \eta_{12}) = \eta_{11}\theta_0 + (1 - \eta_{11})\eta_{12}.$$

Now we see that \hat{Z}_1 is differentiable in θ_0 and

$$I_1(\theta_0, \theta_1) = \mathbf{1}_{\{\theta_1 = \theta_0\}} \in \{0, 1\}.$$

Hence the envelope theorem is applicable and we can write equation (12) for the relaxed problem as

$$\max_{x_1 \geq 0} \left\{ x_1 \left(\theta_1 - (1 - \theta_0) \mathbf{1}_{\{\theta_1 = \theta_0\}} \right) - \frac{x_1^2}{2} \right\}.$$

¹⁹A similar example can be constructed where the canonical representation is well behaved by smoothing the distribution of θ_1 around θ_0 .

The solution to this problem is

$$x_1(\theta_0, \theta_1) = \begin{cases} \theta_1 & \text{if } \theta_1 \neq \theta_0 \\ \max\{0, 2\theta_1 - 1\} & \text{if } \theta_1 = \theta_0 \end{cases}$$

The allocation rule is non-monotone in θ_1 for each $\theta_0 < 1$ and as a result, it cannot be implemented if η_1 is private information. On the other hand, it can be implemented if η_1 is observed by the seller. In this case, the seller knows θ_1 if $\eta_{11} = 0$ and any rule is trivially implementable. If $\eta_{11} = 1$, the solution to the problem is the usual Mussa–Rosen rule. The second part of the allocation rule x_1 above is the optimal scheme for this case. Hence we conclude that x_1 is implementable with publicly observed η_1 .

REFERENCES

- Albright, S. Christian. 1977. "A Bayesian Approach to a Generalized House Selling Problem." *Management Science* 24 (4): 432–40.
- Athey, Susan, and Ilya Segal. 2013. "An Efficient Dynamic Mechanism." *Econometrica* 81 (6): 2463–85.
- Baron, David P., and David Besanko. 1984. "Regulation and Information in a Continuing Relationship." *Information Economics and Policy* 1 (3): 267–302.
- Battaglini, Marco. 2005. "Long-Term Contracting with Markovian Consumers." *American Economic Review* 95 (3): 637–58.
- Battaglini, Marco, and Rohit Lamba. 2018. "Optimal Dynamic Contracting: The First-Order Approach and Beyond." Available on SSRN at 2172414.
- Bergemann, Dirk, Benjamin Brooks, and Stephen Morris. 2017. "Selling to Intermediaries: Optimal Auction Design in a Common Value Model." https://economics.yale.edu/sites/default/files/selling_to_intermediaries.pdf.
- Bergemann, Dirk, Francisco Castro, and Gabriel Weintraub. 2017. "The Scope of Sequential Screening with Ex-Post Participation Constraints." Cowles Foundation Discussion Paper 2078.
- Bergemann, Dirk, and Stephen Morris. 2005. "Robust Mechanism Design." *Econometrica* 73 (6): 1771–813.
- Bergemann, Dirk, and Alessandro Pavan. 2015. "Introduction to Symposium on Dynamic Contracts and Mechanism Design." *Journal of Economic Theory* 159 (B): 679–701.
- Bergemann, Dirk, and Maher Said. 2010. "Dynamic Auctions." In *Wiley Encyclopedia of Operations Research and Management Science*, edited by James J. Cochran, Louis Anthony Cox, Jr., Pinar Keskinocak, Jeffrey P. Kharoufeh, and J. Cole Smith, 1511–22. Hoboken: Wiley.
- Bergemann, Dirk, and Philipp Strack. 2015. "Dynamic

- Revenue Maximization: A Continuous Time Approach." *Journal of Economic Theory* 159: 819–53.
- Bergemann, Dirk, and Philipp Strack. 2019. "Progressive Participation." Unpublished.
- Bergemann, Dirk, and Juuso Välimäki. 2003. "Dynamic Common Agency." *Journal of Economic Theory* 111 (1): 23–48.
- Bergemann, Dirk, and Juuso Välimäki. 2006. "Dynamic Price Competition." *Journal of Economic Theory* 127 (1): 232–63.
- Bergemann, Dirk, and Juuso Välimäki. 2010. "The Dynamic Pivot Mechanism." *Econometrica* 78 (2): 771–90.
- Bergemann, Dirk, and Achim Wambach. 2015. "Sequential Information Disclosure in Auctions." *Journal of Economic Theory* 159 (B): 1074–95.
- Besanko, David. 1985. "Multi-period Contracts between Principal and Agent with Adverse Selection." *Economics Letters* 17 (1–2): 33–37.
- Biais, Bruno, Thomas Mariotti, and Jean-Charles Rochet. 2013. "Dynamic Financial Contracting." In *Advances in Economics and Econometrics*, Vol. 1, edited by Daron Acemoglu, Manuel Arellano, and Eddie Dekel, 125–71. Cambridge: Cambridge University Press.
- Biais, Bruno, Thomas Mariotti, Guillaume Plantin, and Jean-Charles Rochet. 2007. "Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications." *Review of Economic Studies* 74 (2): 345–90.
- Board, Simon. 2007. "Selling Options." *Journal of Economic Theory* 136 (1): 324–40.
- Board, Simon, and Andrzej Skrzypacz. 2016. "Revenue Management with Forward-Looking Buyers." *Journal of Political Economy* 124 (4): 1046–87.
- Boleslavsky, Raphael, and Maher Said. 2013. "Progressive Screening: Long-Term Contracting with a Privately Known Stochastic Process." *Review of Economic Studies* 80 (1): 1–34.
- Börgers, Tilman. 2015. *An Introduction to the Theory of Mechanism Design*. Oxford: Oxford University Press.
- Calzolari, Giacomo, and Alessandro Pavan. 2006. "Monopoly with Resale." *RAND Journal of Economics* 37 (2): 362–75.
- Carroll, Gabriel, and Ilya Segal. 2016. "Robustly Optimal Auctions with Unknown Resale Opportunities." <https://web.stanford.edu/~isegal/resale.pdf>.
- Clementi, Gian Luca, and Hugo A. Hopenhayn. 2006. "A Theory of Financing Constraints and Firm Dynamics." *Quarterly Journal of Economics* 121 (1): 229–65.
- Conlisk, John, Eitan Gerstner, and Joel Sobel. 1984. "Cyclic Pricing by a Durable Goods Monopolist." *Quarterly Journal of Economics* 99 (3): 489–505.
- Courty, Pascal, and Li Hao. 2000. "Sequential Screening." *Review of Economic Studies* 67 (4): 697–717.
- Crémer, Jacques, and Richard P. McLean. 1985. "Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist When Demands Are Interdependent." *Econometrica* 53 (2): 345–61.
- Crémer, Jacques, and Richard P. McLean. 1988. "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions." *Econometrica* 56 (6): 1247–57.
- Dasgupta, Partha, and Eric Maskin. 2000. "Efficient Auctions." *Quarterly Journal of Economics* 115 (2): 341–88.
- Deb, Rahul, and Debasis Mishra. 2014. "Implementation with Contingent Contracts." *Econometrica* 82 (6): 2371–93.
- DellaVigna, Stefano, and Ulrike Malmendier. 2006. "Paying Not to Go to the Gym." *American Economic Review* 96 (3): 694–719.
- DeMarzo, Peter M., and Michael J. Fishman. 2007. "Optimal Long-Term Financial Contracting." *Review of Financial Studies* 20 (6): 2079–128.
- DeMarzo, Peter M., and Yuliy Sannikov. 2006. "Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model." *Journal of Finance* 61 (6): 2681–724.
- Derman, Cyrus, Gerald J. Lieberman, and Sheldon M. Ross. 1972. "A Sequential Stochastic Assignment Problem." *Management Science* 18 (7): 349–55.
- Dolan, Robert J. 1978. "Incentive Mechanisms for Priority Queuing Problems." *Bell Journal of Economics* 9 (2): 421–36.
- Dworczak, Piotr. 2017. "Mechanism Design with Aftermarkets: Cutoff Mechanisms." Available on SSRN at 2859206.
- Eliasz, Kfir, and Ran Spiegler. 2008. "Consumer Optimism and Price Discrimination." *Theoretical Economics* 3: 459–97.
- Eső, Péter, and Balázs Szentés. 2007. "Optimal Information Disclosure in Auctions and the Handicap Auction." *Review of Economic Studies* 74 (3): 705–31.
- Farhi, Emmanuel, and Iván Werning. 2007. "Inequality and Social Discounting." *Journal of Political Economy* 115 (3): 365–402.
- Garrett, Daniel F. 2016. "Intertemporal Price Discrimination: Dynamic Arrivals and Changing Values." *American Economic Review* 106 (11): 3275–99.
- Garrett, Daniel F. 2017. "Dynamic Mechanism Design: Dynamic Arrivals and Changing Values." *Games and Economic Behavior* 104: 595–612.
- Garrett, Daniel F., and Alessandro Pavan. 2012. "Managerial Turnover in a Changing World." *Journal of Political Economy* 120 (5): 879–925.
- Garrett, Daniel, and Alessandro Pavan. 2015. "Dynamic Managerial Compensation: A Variational Approach." *Journal of Economic Theory* 159 (B): 775–818.
- Garrett, Daniel, Alessandro Pavan, and Juuso Toikka. 2018. "Robust Predictions in Dynamic Screening." <http://faculty.wcas.northwestern.edu/~apa522/Robust-dynamics.pdf>.
- Gershkov, Alex, and Benny Moldovanu. 2009a. "Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach." *American Economic Journal: Microeconomics* 1 (2): 168–98.
- Gershkov, Alex, and Benny Moldovanu. 2009b. "Learning about the Future and Dynamic Efficiency." *American Economic Review* 99 (4): 1576–87.

- Gershkov, Alex, and Benny Moldovanu. 2014. *Dynamic Allocation and Pricing: A Mechanism Design Approach*. Cambridge: MIT Press.
- Gershkov, Alex, Benny Moldovanu, and Philipp Strack. 2018. "Revenue-Maximizing Mechanisms with Strategic Customers and Unknown, Markovian Demand." *Management Science* 64 (5): 2031–46.
- Green, Edward J. 1987. "Lending and the Smoothing of Uninsurable Income." In *Contractual Arrangements for Intertemporal Trade*, Vol. 1, edited by Edward C. Prescott and Neil Wallace, 3–25. Minneapolis: University of Minnesota Press.
- Green, Jerry, and Jean-Jacques Laffont 1977. "Characterization of Satisfactory Mechanisms for the Revelation of the Preferences for Public Goods." *Econometrica* 45 (2): 427–38.
- Grubb, Michael D. 2009. "Selling to Overconfident Consumers." *American Economic Review* 99 (5): 1770–807.
- Grubb, Michael D., and Matthew Osborne. Forthcoming. "Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock." *American Economic Review*.
- Harris, Milton, and Bengt Holmström. 1982. "A Theory of Wage Dynamics." *Review of Economic Studies* 49 (3): 315–33.
- Holmström, Bengt. 1979. "Groves' Scheme on Restricted Domains." *Econometrica* 47 (5): 1137–44.
- Kakade, Sham M., Ilan Lobel, and Hamid Nazerzadeh. 2013. "Optimal Dynamic Mechanism Design and the Virtual-Pivot Mechanism." *Operations Research* 61 (4): 837–54.
- Krähmer, Daniel, and Roland Strausz. 2015. "Optimal Sales Contracts with Withdrawal Rights." *Review of Economic Studies* 82 (2): 762–90.
- Kuribko, Nataliya, Tracy Lewis, Fang Liu, and Jing-Sheng Jeannette Song. 2017. "Long Term Partnership for Achieving Efficient Capacity Allocation." Available on SSRN at 2518046.
- Li, Hao, and Xianwen Shi. 2017. "Discriminatory Information Disclosure." *American Economic Review* 107 (11): 3363–85.
- Liu, Heng. 2013. "Efficient Dynamic Mechanisms in Environments with Interdependent Valuations." <http://www.gtcenter.org/Archive/2013/Conf/Liu1630.pdf>.
- Mezzetti, Claudio. 2004. "Mechanism Design with Interdependent Valuations: Efficiency." *Econometrica* 72 (5): 1617–26.
- Mezzetti, Claudio. 2007. "Mechanism Design with Interdependent Valuations: Surplus Extraction." *Economic Theory* 31 (3): 473–88.
- Mierendorff, Konrad. 2013. "The Dynamic Vickrey Auction." *Games and Economic Behavior* 82: 192–204.
- Mierendorff, Konrad. 2016. "Optimal Dynamic Mechanism Design with Deadlines." *Journal of Economic Theory* 161: 190–222.
- Milgrom, Paul, and Ilya Segal. 2002. "Envelope Theorems for Arbitrary Choice Sets." *Econometrica* 70 (2): 583–601.
- Mirroknii, Vahab, Renato Paes Leme, Pingzhong Tang, and Song Zuo. 2018. "Non-clairvoyant Dynamic Mechanism Design." <https://ai.google/research/pubs/pub47744>.
- Mookherjee, Dilip, and Stefan Reichelstein. 1992. "Dominant Strategy Implementation of Bayesian Incentive Compatible Allocation Rules." *Journal of Economic Theory* 56 (2): 378–99.
- Moulin, H. 1986. "Characterization of the Pivotal Mechanism." *Journal of Public Economics* 31 (1): 53–78.
- Myerson, Roger B. 1986. "Multistage Games with Communication." *Econometrica* 54 (2): 323–58.
- Papadimitriou, Christos, George Pierrakos, Christos-Alexandros Psomas, and Aviad Rubinfeld. 2016. "On the Complexity of Dynamic Mechanism Design." In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, edited by Robert Krauthgamer, 1458–75. Philadelphia: Society for Industrial and Applied Mathematics.
- Pavan, Alessandro. 2017. "Dynamic Mechanism Design: Robustness and Endogenous Types." In *Advances in Economics and Econometrics: Eleventh World Congress*, edited by Bo Honoré, Ariel Pakes, Monika Piazzesi, and Larry Samuelson, 1–62. Cambridge: Cambridge University Press.
- Pavan, Alessandro, Ilya Segal, and Juuso Toikka. 2014. "Dynamic Mechanism Design: A Myersonian Approach." *Econometrica* 82 (2): 601–53.
- Riordan, Michael, and David E. M. Sappington. 1987. "Information, Incentives, and Organizational Mode." *Quarterly Journal of Economics* 102 (2): 243–63.
- Rogerson, William P. 1985. "The First-Order Approach to Principal-Agent Problems." *Econometrica* 53 (6): 1357–67.
- Skrzypacz, Andrzej, and Juuso Toikka. 2015. "Mechanisms for Repeated Trade." *American Economic Journal: Microeconomics* 7 (4): 252–93.
- Sugaya, Takuo, and Alexander Wolitzky. 2017. "Revelation Principles in Multistage Games." <https://pdfs.semanticscholar.org/05bc/d4d74bb3be8c2612549c216d605a26bca2e0.pdf>.
- Talluri, Kalyan T., and Garrett J. Van Ryzin. 2004. *The Theory and Practice of Revenue Management*. New York: Springer.
- Thomas, Jonathan, and Tim Worrall. 1990. "Income Fluctuation and Asymmetric Information: An Example of a Repeated Principal-Agent Problem." *Journal of Economic Theory* 51 (2): 367–90.