

**INTRODUCTION TO JET SYMPOSIUM ISSUE
ON “DYNAMIC CONTRACTS AND MECHANISM DESIGN”**

By

Dirk Bergemann and Alessandro Pavan

August 2015

COWLES FOUNDATION DISCUSSION PAPER NO. 2016



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

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

Introduction to JET Symposium Issue on "Dynamic Contracts and Mechanism Design"

Dirk Bergemann* Alessandro Pavan†

July 31, 2015

Abstract

The Introduction to the Symposium Issue on "Dynamic Contract and Mechanism Design" of the Journal of Economic Theory provides an overview of the dynamic mechanism design literature. We then introduce the papers that are contained in the Symposium issue and finally conclude by discussing avenues for future research. Several of the papers contained in the Symposium issue were presented at the Economic Theory Workshop of the Cowles Foundation for Research in Economics at Yale University in June 2013.

KEYWORDS: Dynamic Mechanism Design, Dynamic Contracts.

JEL CLASSIFICATION: D44, D82, D83.

*Department of Economics, Yale University, New Haven, CT 06511, dirk.bergemann@yale.edu

†Department of Economics, Northwestern University, Evanston, IL 60208, alepavan@northwestern.edu

1 Introduction

Mechanism design has proved to be a powerful tool to examine a variety of phenomena, including auctions, nonlinear pricing, regulation, taxation, political economy, the provision of public goods, and the design of organizations.¹ A recent monograph, Börgers (2015), provides an excellent and comprehensive overview of the theory of mechanism design.

Most of the early path-breaking work assumed that information is dispersed but given once and for all at the outset of the game. It also confined attention to settings with a single allocative decision to be taken. For many problems of interest, though, information arrives gradually over time and there is a stream of decisions to be made. Think, for example, of the allocation of private or public goods to agents whose valuations evolve over time as the result of experimentation, or the design of multi-period procurement auctions when firms' costs are endogenous and are the firms' private information.

As a result, the last fifteen years have witnessed significant interest in extending the theory of mechanism design to dynamic environments. This Symposium advances our understanding of dynamic contracts and mechanism design by bringing together recent contributions that cover a wide range of issues at the research frontier of the dynamic mechanism design literature. Several of these papers were presented at the Economic Theory Workshop of the Cowles Foundation for Research in Economics at Yale University in June 2013.

In this Introduction, we first provide an overview of the dynamic mechanism design literature. We then introduce the various papers, and finally conclude by discussing avenues for future research.

¹For earlier contributions to auctions, see Myerson (1981) and Riley and Samuelson (1981); to nonlinear pricing, see Mussa and Rosen (1978), and Wilson (1993); to regulation, see Baron and Myerson (1982) and Laffont and Tirole (1986); to taxation, see Mirrlees (1971); to political economy, see Dasgupta, Hammond, and Maskin (1979), and Acemoglu, Golosov, and Tsyvinski (2011); to the provision of public goods, see Vickrey (1961), Clarke (1971), Groves (1973), and Green and Laffont (1979); to the design of organizations, see Cremer (1995).

2 A Brief Review of the Dynamic Mechanism Design Literature

A stream of the literature investigates how to implement dynamically efficient allocations in settings in which the agents' types change over time, thus extending the Vickrey–Clarke–Groves (VCG) and d'Aspremont–Gérard-Varet (AGV) results from static to dynamic settings (see, for example, Bergemann and Välimäki (2010), Athey and Segal (2013) and the references therein). The dynamic pivot mechanism and the dynamic AGV mechanism are shown to preserve most of the properties of their static analogs: the dynamic pivot mechanism guarantees that, in each period, all agents receive their expected marginal contribution to social welfare, thus guaranteeing participation in all periods, whereas the dynamic AGV mechanism achieves budget balance in all periods.

Another stream of the literature investigates the design of revenue-maximizing mechanisms in dynamic settings. Earlier contributions include Baron and Besanko (1984), Besanko (1985), and Riordan and Sappington (1987); for more recent contributions, see Courty and Li (2000), Battaglini (2005), Esó and Szentes (2007), Board (2007), and Kakade, Lobel, and Nazerzadeh (2013).

The approach typically followed in the design of optimal mechanisms consists in first identifying necessary conditions for incentive compatibility that permit one to express transfers as a function of the allocation rule and express the principal's objective as dynamic virtual surplus. The second step then consists in optimizing dynamic virtual surplus across all possible allocation rules, including those that are potentially not incentive compatible. The third and final step consists in verifying that the allocation rule that solves the relaxed program (along with a transfer rule that guarantees that all the local constraints hold) constitute a fully incentive-compatible and individually-rational mechanism. This last step typically involves identifying appropriate primitive conditions that guarantee that the allocation rule that solves the relaxed program is sufficiently monotone in an appropriate dynamic sense.

The approach described above, which traces back to Myerson (1981), Guesnerie and Laffont (1984), and Maskin and Riley (1984), has been recently extended to dynamic problems by Pavan, Segal, and Toikka (2014). They consider a general dynamic model with a continuum of types, in which agents receive private information over time, and decisions are made in multiple periods over an arbitrary time horizon. The model allows for serial correlation of the agents' information and

for the dependence of this information on past allocations.

The cornerstone in Pavan, Segal, and Toikka (2014) is a dynamic envelope theorem that yields a formula for the evolution of the equilibrium payoffs. As in static models (see Mirrlees (1971) or Myerson (1981)), this formula characterizes local incentive-compatibility constraints. It combines the usual direct effect of a change in the current type on the agent’s utility with novel indirect effects stemming from the induced change in the distribution of the agent’s future types. The stochastic component of the latter is summarized by *impulse response functions* that describe how a change in the agent’s current type propagates through the entire type process.

A second key contribution of Pavan, Segal, and Toikka (2014) is in showing how, in Markov environments, the aforementioned dynamic envelope formula, paired with an appropriate dynamic *integral monotonicity* condition, identifies global incentive compatibility.² The integral monotonicity condition is weaker than the familiar notion of *strong monotonicity* typically considered in the literature, which requires that each agent’s allocation be increasing in each of his current and past reports in every period. Integral monotonicity requires, instead, that allocations be monotone in types “on average” where the average is both across time and future types.

As in static settings, the first-order approach described above yields an implementable allocation rule only under fairly stringent conditions. An important question for the dynamic mechanism design literature is thus the extent to which the predictions identified under such an approach extend to environments where global incentive constraints bind. The paper by Garrett and Pavan (2015) in this Symposium is an attempt in this direction.³

Another stream of the literature considers both efficient and profit-maximizing mechanisms in settings where the agents’ private information is static, but where interesting dynamics originate by agents or objects arriving stochastically over time, as in Gershkov and Moldovanu (2009a), (2009b), (2012) and their recent monograph, Gershkov and Moldovanu (2014), as well as Board and Skrzypacz (2015) and Gershkov, Moldovanu, and Strack (2014); see Bergemann and Said (2011) for

²In static settings, the connection between implementability and monotonicity conditions analogous to integral monotonicity was first noticed by Rochet (1987) and more recently in Carbajal and Ely (2013) and Berger, Müller, and Naeemi (2010). In a dynamic setting, Rahman (2011) characterizes implementable rules based on a monotonicity condition similar to Rochet’s cyclical monotonicity.

³Garrett, Pavan, and Toikka (2015) also use variational arguments which they apply directly to the full program to identify certain robust predictions of optimal dynamic screening models.

an overview of this recent literature). In most of the contributions to this literature, the agents arrive stochastically over time but perfectly learn their values upon arrival, as for example in Said (2011), (2012). In contrast, Garrett (2014), Hinnosaar (2015), and Ely, Garrett, and Hinnosaar (2015) consider models in which agents gradually arrive *and* learn their values over time. A maintained assumption is that the time at which each agent learns his valuation is independent of the realized valuations. The paper by Akan, Ata, and Dana (2015) in the present Symposium studies the implications of relaxing this independence assumption.

The applications in most of the work discussed above are to auctions, revenue management, nonlinear pricing, and regulation. Dynamic mechanism design has also been used to study the design of insurance and optimal taxation in dynamic economies. Earlier contributions include Green (1987), Atkeson and Lucas (1992), and Fernandes and Phelan (2000). For more recent contributions, see Kocherlakota (2005), Albanesi and Sleet (2006), Farhi and Werning (2013), Kapicka (2013), Stantcheva (2014) and Golosov, Tsyvinski, and Troshkin (2015). This literature considers economies in which the agents' private information evolves stochastically over time, but is exogenous to the agents' decisions. In contrast, Makris and Pavan (2015) consider a model of dynamic taxation under learning by doing, where the agents' private information is endogenous and depends on the amount of effort exerted in previous periods.⁴

Related to the optimal taxation literature is the literature on dynamic managerial compensation. Most of this literature is concerned with the provision of incentives for effort in a pure moral hazard setting (see, for example, Sannikov (2013) for an overview of the continuous-time contracting literature, and Board (2011) for a model of relational contracting). The part of this literature that is mostly related to dynamic mechanism design is the one that assumes that the manager observes privately a payoff or productivity shock prior to committing his effort (as in the taxation literature); see, for example, Edmans and Gabaix (2011), Edmans, Gabaix, Sadzik, and Sannikov (2012), and Garrett and Pavan (2012).

The analysis in most of the dynamic mechanism design literature is in discrete time. A notable earlier exception is Williams (2011). The paper by Bergemann and Strack (2015a) in this Sympo-

⁴The information is also determined endogenously by the mechanism in Kakade, Lobel, and Nazerzadeh (2013), and Fershtman and Pavan (2015). These papers consider intertemporal models of experimentation in which the agents revise their beliefs through experience and consumption. Similar issues arise in the application to licensing in Bergemann and Välimäki (2010) and in the multi-armed bandit-auction in Pavan, Segal, and Toikka (2014).

sium provides an illustration of how to extend mechanism design to continuous-time models and of how the latter may facilitate novel and sharper predictions. In their leading example of repeat sales of a good or service, they establish that commonly observed contract features such as flat rates, free consumption units, and two-part tariffs are natural parts of optimal contracts. The contributions by Strulovici and Szydlowski (2015) and Williams (2015) in this Symposium develop some of the associated optimization results and derive explicit solutions to a class of dynamic principal-agent problems, respectively.

All the above papers assume that the designer can commit to her mechanism, in which case the dynamics are driven either by changes in the agents' information or by the stochastic arrival of goods and agents over time. A different stream of the literature studies dynamic contracting in settings in which the agents' private information is static but where interesting dynamics originate via the principal's lack of commitment (for earlier contributions, see, for example, Laffont and Tirole (1988), and Hart and Tirole (1988); for more recent contributions, see Skreta (2006) and the references therein).⁵ Most works on limited commitment consider settings with a single agent. The paper by Skreta (2015) in the present Symposium, instead, considers an auction environment with multiple bidders in which the designer lacks the commitment not to redesign her mechanism in case the good remains unassigned.

A different form of limited commitment is considered in the paper by Deb and Said (2015) also in the present Symposium. In that paper the seller can commit to the dynamic contract she offers to each agent, but cannot commit to the contracts she will offer to future cohorts of arriving agents. Partial commitment is also the focus of the paper by Miao and Zhang (2015) in the present Symposium. They consider a dynamic insurance environment with symmetric information, in which both the principal and the agent can walk away from the relationship after observing the evolution of the agent's income process.

Dynamic mechanism design is also related to the literature on persuasion and optimal information disclosure (see Kamenica and Gentzkow (2011) and the references therein). The canonical persuasion model assumes that the designer (the sender) can choose the information structure of the agent (the receiver). Interesting effects emerge when one combines persuasion with screening. Calzolari and Pavan (2006a), (2006b), for example, consider models in which a principal first screens

⁵See also Battaglini (2007) and Strulovici (2010) for models of limited commitment with changing types.

the private information of one or multiple agents and then passes some of this information to other agents, or other principals. Bergemann and Pesendorfer (2007) and Esó and Szentes (2007) in turn study the design of optimal information structures in auctions. The paper by Bergemann and Wambach (2015) in the present Symposium contributes to this literature by allowing for sequential disclosure policies.

Most of the papers in the literature with time-varying private information assume that the agents report their private information over time. In other words, they assume that the dynamic allocations are implemented with multiple rounds of communication. The paper by Kruse and Strack (2015) in the present Symposium, instead, considers a family of stopping problems where the optimal allocations can be sustained either without communication or with a single round of initial communication.

The last paper in the Symposium by Bognar, Börgers, and Meyer-ter-Vehn (2015) considers an interesting application of dynamic mechanism design to sequential voting. It is related to the literature that studies information acquisition in mechanism design (see Bergemann and Välimäki (2002), (2006) and the references therein for earlier contributions, and Gershkov and Szentes (2009) and Krähmer and Strausz (2011) for recent developments).

3 Overview of the Symposium Contributions

The first paper in the Symposium by **Kruse and Strack (2015)** analyzes an important class of dynamic allocation problems with private information, the class of optimal stopping problems. In a discrete-time setting, a single agent privately observes a stochastic process. Based on the private observation of the process, the agent then decides when to stop the process, and the decision to stop is irreversible. The agent's and the principal's gross utilities depend on the state of the stochastic process before the time, at the time, and after the time of stopping. The principal can affect the decision of the agent by offering a transfer that is contingent on the agent's stopping decision.

The first objective of the paper is to describe the set of implementable policies. In this environment, a direct mechanism would require the agent to report at every period the current state of the process. Because the optimal reporting strategies may not be Markovian, the resulting optimization

of the agent, and in turn of the principal, may be hard to solve. Instead, Kruse and Strack (2015) restrict attention to mechanisms that can be implemented without communication. In other words, the transfer that the principal offers to the agent depends only on the time the process is stopped. The paper then derives a dynamic version of the familiar single-crossing condition under which all cut-off rules can be implemented without communication. A cut-off rule is one in which stopping occurs in period t if and only if the agent's type (described by an increasing function of the process) is above a given cut-off. The converse result is also shown to hold. Namely, if an allocation rule is implementable without communication, then it has to be a finite cut-off rule.

While the focus on mechanisms without communication is clearly with loss of generality with respect to the larger class of direct mechanism, Kruse and Strack (2015) show that such mechanisms are often optimal with respect to the broader class of unrestricted mechanisms. For example, the socially efficient allocation in an economy in which the agent's and the principal's utilities depend only on the flow value of the process can always be implemented without communication. The paper also shows that, when some communication is needed to attain the principal's supremum payoff, it can often be limited to the first period (i.e., the principal offers a menu of stopping rules, indexed by the initial state of the process). For example, in the case of additive random walks, the maximal revenue for the principal can be attained by restricting communication to the initial period. This result mirrors an earlier one by Board (2007) that analyzes the revenue-maximizing sale of real options under private information.

Kruse and Strack (2015) explicitly construct the transfers that implement any given cut-off rule. In order to induce stopping at a given cut-off value, the agent needs to be compensated for the forgone option of stopping *one*, or possibly *many*, periods later. By contrast, the single-crossing condition in Kruse and Strack (2015) is defined in terms of the marginal value of waiting one extra period. Surprisingly, the paper shows that the transfers implementing a cut-off rule can be constructed as the sum of all consecutive expected marginal option values. However, the expectation has to be taken with respect to a modified version of the stochastic process. Namely, the modified process is a constrained version of the original process in the sense that the process is bounded precisely by the sequence of cut-off values that is to be implemented. In other words, any *upward increment* in the original process beyond the cut-off value is simply replaced by a *zero increment*, and as soon as the process drops below the cut-off value it becomes unconstrained

again. Thus, the agent is compensated for the entire option value across all future periods, but the expectation is taken with respect of the constrained version of the stochastic process, which reflects the implementation of the cut-off policy.

The second paper in the Symposium by **Akan, Ata, and Dana (2015)** studies a revenue-management problem in which a seller faces a population of buyers who are privately informed about the time at which they will learn their valuations. The time at which the agents learn is also correlated with the agents' valuations and this gives rise to interesting novel effects. In particular, the seller can screen on both the size of the refund and on the time the refund option expires.

Intuitively, screening on when the refund option expires may allow the seller to slack the agents' "downward" local incentive compatibility constraints. This is because an agent who claims to learn earlier than she actually does forgoes the option value of returning the good in case the true valuation turns out to be low.

The paper then shows that, in some cases, screening along the time dimension may permit the seller to extract the full surplus, despite the fact that the agents possess private information at the time of contracting. In particular, the paper shows that full surplus extraction is possible when those consumers who learn later have a distribution of valuations that is a mean-preserving spread of the distribution of valuations of those consumers who learn earlier. In this case, the option value of returning the good is the highest for those consumers who learn later. The seller can then charge a "premium" for the option to return the good that is increasing in the deadline for the refund option and use a menu of such refunds to extract the entire surplus.

When, instead, the expected gains from trade are decreasing in the time the agents learn their valuations, screening on the time dimension has no effect on profits. In this case, if the seller were to offer the menu of complete-information contracts, all consumers would claim to learn late and enjoy an informational rent. In fact, this case resembles the more familiar screening setting in which all consumers learn their values at the same time and future values are increasing in the agents' early private information in a first-order-stochastic-dominance sense (see, e.g., Courty and Li (2000)).

More generally, the paper shows that it is optimal to screen both on the size of the refund and on the time the refund option expires when (a) the expected gains from trade are highest for consumers who learn latest, and (b) the value of the return option is not high enough to deter consumers who learn late from imitating consumers who learn early.

Interestingly, the paper shows that when (i) the distribution of valuations of later-learning consumers first-order-stochastically dominates that of early-learning consumers and (ii) the complete information allocations are not incentive compatible, the menu of optimal contracts can feature both downward and upward distortions. In particular, the seller offers the lower-valuation, early-learning consumers contracts with a higher refund price (thus inducing a downward distortion in consumption) and the higher-valuation, later-learning consumers contracts with a lower refund price (thus inducing upward distortion in consumption). The distortion for the early-learning consumers is the familiar one (it deters the later-learning, higher-valuations consumers from mimicking). However, contrary to standard screening models, only the highest type is indifferent between reporting truthfully and imitating those consumers who learn earlier. Furthermore, the highest type never uses the refund option when she imitates an earlier-learning consumer because she does not know yet her final valuation at the time the refund option expires. In turn, the lower return price (the upward distortion) in the contracts for the late-learning consumers is to deter mimicking by the early-learning consumers. This is a novel binding “upward” local incentive-compatibility constraint that is absent in standard models where all agents learn at the same time.

Many vendors (airlines, hotels, event organizers) offer menus of contracts with sophisticated refund policies. The paper provides a useful framework for thinking about the relevant trade-offs and offers valuable insights on the practices followed by some of these vendors. It brings together results from the dynamic mechanism design literature and the operations management literature on pricing and capacity controls in stochastic environments. As most of the papers in this Symposium (see Skreta (2015), and Deb and Said (2015) for exceptions), Akan, Ata, and Dana (2015) assume that the seller can perfectly commit to her mechanism. Extending the analysis to settings in which the seller has imperfect commitment is likely to bring new insights and make the predictions further aligned with actual pricing practices.

The third paper in the Symposium by **Garrett and Pavan (2015)** considers a dynamic contracting environment in which a manager’s ability to generate profits for a firm changes stochastically with time. Both the manager’s ability and his effort are the manager’s private information. Importantly, the manager possesses private information prior to contracting with the firm. The purpose of the analysis is to examine the dynamics of incentives under optimal contracts.

In this environment, asking a manager to exert more effort is costly for three reasons. First,

higher effort is costly for the manager and hence must be compensated. Second, asking higher effort of a manager of a given productivity requires increasing the compensation to all managers with higher productivity; this extra compensation is required to discourage these more productive managers from mimicking the less productive ones. Lastly, higher effort requires pay to be more sensitive to performance. When the manager is risk averse, this increase in volatility requires higher compensation by the firm.

The paper investigates the implications of the above effects both for the dynamics of effort and for the distortions in the provision of incentives due to asymmetric information. As in the New Dynamic Public Finance, in the presence of wealth effects (that is, beyond the quasilinear case), distortions are best measured by the “wedge” between the marginal cash flows generated by higher effort and the marginal compensation that must be paid to the managers to keep their utility constant. The dynamics of wedges provide information on how the firm optimally distorts both effort and compensation intertemporally to reduce the managers’ information rents.

Contrary to most of the dynamic mechanism design literature, the paper identifies certain properties of optimal contracts by applying variational arguments directly to the firm’s “full problem”. It identifies perturbations that preserve participation and incentive-compatibility constraints and then uses such perturbations to identify properties of optimal contracts. This approach does not permit one to fully characterize how effort and compensation respond to all possible contingencies. However, it does permit one to identify how, on average, effort and incentives evolve over time under fully optimal contracts.

The advantage of this variational approach is that it permits one to bypass some of the difficulties in the dynamic mechanism design literature. As mentioned above, the typical approach involves imposing only a restricted set of “local” incentive constraints, then solving a “relaxed problem,” and finally identifying restrictions on the primitive environment that guarantee that the solution to the relaxed problem satisfies the remaining constraints. When validated, this “first-order” approach has the advantage of yielding ex-post predictions about effort and compensation that depend on the realized productivity history. In contrast, the variational approach in Garrett and Pavan (2015) yields only ex-ante predictions that hold by averaging over productivity histories.

In terms of predictions, the paper shows that, when the managers are risk neutral, the firm typically distorts downward (relative to the first best) the level of effort asked of those managers

whose initial productivity is low. Importantly, this result is shown to extend to settings in which the first-order approach fails, provided that effort is bounded away from zero from below (the paper provides primitive conditions for this to be the case). It is also shown that, whenever (a) on average, first-period effort is distorted downward relative to the first-best level, and (b) the effect of the initial productivity on future productivity declines with time, the firm asks, on average, for higher effort later in the relationship. The result follows from the fact that, when productivity is less than fully persistent, the benefit of distorting the effort of those managers whose initial productivity is low so as to reduce the compensation paid to those managers whose initial productivity is high is greatest early in the relationship.

The most interesting results pertain to the case in which the managers are risk averse. To reduce the volatility of future compensation the firm further distorts future effort and compensation away from their efficient levels. Whether distortions increase or decrease, on average, over time then depends on the degrees of managerial risk-aversion and productivity persistence. For low degrees of risk aversion and low degrees of productivity persistence, the dynamics of distortions parallel those in the risk neutral case (that is, distortions decrease, on average, over time). When, instead, shocks to productivity are sufficiently persistent, as in the case of a random walk, then, for any degree of risk aversion, distortions increase, on average, over time.

The variational approach developed in Garrett and Pavan (2015) could turn useful also in other contractual environments. For example, in ongoing work, Garrett, Pavan, and Toikka (2015) use a similar approach to identify long-run properties of optimal screening contracts in a Mussa–Rosen environment of nonlinear pricing. They show how convergence to efficiency depends on properties of the process such as ergodicity and first-order stochastic dominance.

The fourth paper by **Bergemann and Strack (2015a)** considers the classic problem of dynamic revenue maximization under private information. In contrast to much of the received literature in dynamic mechanism design, the analysis is in continuous rather than discrete time. The private information of the agent is assumed to follow a stochastic process with a Brownian noise term. This includes, among others, the arithmetic and the geometric Brownian motion, but also the mean-reverting Ornstein–Uhlenbeck process, or certain types of Bayesian learning. The focus of the analysis is on time-separable allocation problems in which the feasibility of the current set of allocations is unaffected by past decisions. This class includes the repeated sale of goods or services,

but does exclude stopping problems, such as the single sale of a durable good in which past allocative decisions restrict the set of feasible contemporaneous allocations.

The specific insights of the continuous-time setting emerge after the necessary conditions for the optimal mechanism are established. Specifically, while the necessary conditions of the continuous-time setting mirror almost exactly those of the discrete-time setting (see Esó and Szentes (2007) and Pavan, Segal, and Toikka (2014)), the sufficient conditions for the time-separable problem in continuous time become very transparent. In particular, using the familiar tools from stochastic calculus, Bergemann and Strack (2015a) can explicitly derive the solution to the dynamic mechanism design problem. In the leading case of the repeat sales of a good or service, Bergemann and Strack (2015a) identify conditions under which commonly observed features such as flat rates, free consumption units, or two-part tariffs arise as part of the optimal mechanism.

An early and important insight in the dynamic mechanism design literature (see Baron and Besanko (1984)) is that, relative to the static problem, the virtual utility of the agent has to be augmented by a term that represents the informational impact of the agent's initial private information at time zero on the future types. Pavan, Segal, and Toikka (2014) show how this informational term can be written for general stochastic processes as a function of orthogonalized signals – the impulse responses of the stochastic process. In the continuous-time Brownian motion setting of Bergemann and Strack (2015a), the incremental information is by construction orthogonal to the past information. Moreover, the continuous time limit of the impulse response function is simply the stochastic flow of the continuous time process, an expression that is frequently very compact and described in terms of the primitives of the stochastic process, time, mean, and variance. Thus, the characterization of the optimal mechanism and the nature of the allocative distortion is often simpler in the continuous-time setting relative to its discrete-time analog.

In Bergemann and Strack (2015a), the agent's private information at every time $t > 0$ consists of the state of the stochastic process. The agent's private information at time $t = 0$, instead, is allowed to be different from the initial state of the process. For example, it can be the drift, or the variance, of the stochastic process or, in the case of the mean-reverting process, the speed of the mean-reversion. In those instances, the corresponding initial state of the process is public information. Thus, at every point in time, the agent only receives a one-dimensional private signal and the analysis of the incentive constraints remains confined to a sequence of one-dimensional incentive constraints.

Yet, by allowing for a richer class of initial private signals, Bergemann and Strack (2015a) can analyze the distortionary influence of various forms of initial private information on subsequent allocative decisions. In the Markovian framework considered in much of the earlier literature, the influence of the initial state on the future state vanishes under mild ergodicity conditions. In contrast, Bergemann and Strack (2015a) show that, in many instances, the distortionary influences may remain constant (as in the arithmetic Brownian motion when the initial value of the process is privately known), increase over time (as in the geometric Brownian motion when the drift is privately known) or vary randomly (as in the arithmetic Brownian motion when the variance is privately known). In recent contributions in discrete time settings, Boleslavsky and Said (2013) and Skrzypacz and Toikka (2015) also observe that if the initial private information may pertain to a parameter of the stochastic process itself, then the intertemporal distortion may not disappear over time. Bergemann and Strack (2015a) provide general necessary and sufficient conditions regarding the long-run behavior of distortions in terms of conditions on the convergence behavior of the underlying stochastic process.

The fifth paper in the Symposium by **Skreta (2015)** considers the problem of a seller designing revenue-maximizing mechanisms in a setting in which she cannot commit not to propose a new mechanism if the one previously chosen failed to allocate the object. The environment is similar to the one in Myerson (1981). A risk-neutral seller who owns a single object faces a fixed population of potential risk-neutral buyers whose valuations are private, independently distributed across buyers, and constant over time. At the beginning of each period, the seller proposes a new mechanism. If the object is sold, the game ends; otherwise, in the next period the seller offers a new mechanism. The game ends after an exogenously fixed number of periods.

The key result is that, despite the lack of commitment, the optimal mechanism takes the same form as in Myerson (1981). In particular, when buyers are ex-ante identical, in each period the seller uses first-price (or second-price) auctions with optimally chosen reserve prices. When, instead, buyers are asymmetric (in the sense that their valuations are drawn from different distributions), in each period the sequentially-optimal mechanism assigns the good to the buyer with the highest virtual valuation, provided it is above a buyer-specific reserve price. Importantly, the virtual valuations, and the corresponding reserve prices, are computed by truncating the distributions using the agents' bidding strategies in previous periods. Under the optimal mechanism, a positive-measure

set of types pools by bidding under the reserve price. Finally, the paper shows that the revenue loss due to the lack of commitment is highest for intermediate values of the discount factor and when the number of buyers is small.

The paper is related to the durable-good-monopolist literature (e.g., Stokey (1981), Bulow (1982), Gul, Sonnenschein, and Wilson (1986)). There are two key differences with respect to this literature. First, the seller has limited capacity and hence uses auctions as opposed to posted prices. Second, the model in Skreta (2015) assumes a finite horizon, while most of these earlier works consider a setting with infinitely many periods. The paper is also related to McAfee and Vincent (1997) who study an auction environment in which the seller behaves sequentially rationally, but where she is restricted to choosing reservation prices in an exogenously given auction format (see also Burguet and Sákovic (1996), Caillaud and Mezzetti (2004) and Liu, Mierendorff, and Shi (2015) for alternative models of sequential auctions in which the trading protocol is exogenous).

Methodologically, this is the first paper that solves for the optimal mechanism under limited commitment in a multi-agent environment allowing for a continuum of valuations and for the possibility that the seller controls what agents observe—the transparency of the mechanisms. As is well known, mechanism design under limited commitment is difficult, for one cannot restrict the agents to truthfully reporting their types (see, among others, Freixas, Guesnerie, and Tirole (1985), and Laffont and Tirole (1988)). A remarkable result, first noticed in Kumar (1985) and in Laffont and Tirole (1990) and then subsequently generalized in Bester and Strausz (2001), is that, with a single agent and finitely-many types, the maximal payoff for the principal can be attained by offering mechanisms in which the message space has the same cardinality as the type space, and by inducing the agent to report his true type with strictly positive probability. Unfortunately, as shown in Bester and Strausz (2000), in general, this result does not extend to settings with multiple agents.

In the multi-agent environment in Skreta (2006), optimal allocations can be identified by restricting the seller to offering in period one a dynamic mechanism that is PBE-feasible. Suppose there are only two trading periods. The mechanism specifies a period-one allocation and payment rule. In case the good remains unassigned, the mechanism then specifies a period-two allocation and payment rule. In both periods, each agent’s message space can be restricted to his type space. As in Bester and Strausz (2001), agents are induced to lie with positive probability. The requirement that

the dynamic mechanism be PBE-feasible then amounts to requiring that (a) the proposed mechanism admit a PBE in the game among the agents only and (b) given the PBE played by the agents, the seller not have an incentive in period two to renege and offer a different incentive-compatible direct revelation mechanism. This step then permits the author to reduce the problem of designing a sequence of profit-maximizing auctions to a constrained optimization problem, as in settings with full commitment. The optimization is, however, significantly harder for the following reasons: (i) the agents' optimal reporting strategies cannot be assumed to be truthful, (ii) the allocations must also satisfy the seller's incentive-compatibility constraints (which in turn depend on the agents' reporting strategies), (iii) the design must also specify what the agents observe at the end of the first period (i.e., a disclosure policy), and (iv) in case the seller deviates in period two, she faces an informed principal problem (due to the fact that she may have observed more information than the agents in the first period). Much of this complexity is absent in single-agent settings, such as those considered in Skreta (2006). Remarkably, Skreta (2015) shows that, at the optimum, the seller offers a fully transparent mechanism and induces all types below a given threshold to pool on the same report.

The paper is also related to the literature that studies the design of optimal disclosure policies in settings in which the agents possess private information. Calzolari and Pavan (2006a) show how to design optimal disclosure policies in an auction setting in which the seller cannot prevent resale, whereas Calzolari and Pavan (2006b) show how to design optimal policies in settings in which multiple principals sequentially contract with the same agent. In turn, Bergemann and Pesendorfer (2007), Esó and Szentes (2007), Li and Shi (2015), and Bergemann and Wambach (2015) study optimal disclosure policies in auctions, but where the information disclosed to the bidders is exogenous to the bidders, as opposed to being determined by other bidders' reporting strategies.

The sixth paper in the Symposium by **Deb and Said (2015)** contributes to an open and important question in dynamic mechanism design. It reconsiders the seminal model of sequential screening of Courty and Li (2000), but allows for the arrival of new buyers over time. In Courty and Li (2000), as well as in much of the subsequent literature, the seller can observe the arrival of each buyer and can commit to exclude from future trade any buyer who refuses to contract upon arrival. As a result, a buyer cannot strategically delay his interaction with the seller. This restriction is

definitely at odds with many markets. Consider, for the example, the sale of cell phone contracts with various amounts of free minutes, various data packages, and different commitment times. The seller clearly anticipates she will continue selling new contracts to newly arriving buyers but, to the extent that the buyers are anonymous, she cannot distinguish between a buyer who chooses to delay his interaction with the seller and one arriving late. In other words, arrivals are rarely observable in most markets of interest.

Deb and Said (2015) capture the key ingredients of these situations by considering a two-period model in which the seller can offer long-term contracts to each buyer arriving in period one but cannot commit in the first period to the contracts she will offer to those buyers arriving in the second period. The long-term contracts offered to the first period buyers take the form of option contracts and are meant to provide the early buyers with some flexibility in the way they will respond to the arrival of subsequent private information. Importantly, the seller cannot distinguish between a buyer arriving in the second period and one arriving in the first period but delaying contracting till the second period.

The possibility of timing the contracting date creates an endogenous outside option for the early-arriving buyers. This option in turn restricts the seller's ability to extract surplus from the early-arriving buyers. As a result, the seller has an ex-ante incentive to manage the future contracts in order to make the outside option less valuable to the early buyers. Because the seller is constrained to behave sequentially rationally in the second period, she may then want to manipulate the composition of the second period demand so as to credibly offer less generous terms to those buyers arriving late.

The contribution in Deb and Said (2015) is in showing how a careful demand management that naturally arises under limited commitment can change the properties of optimal contracts. A key result is that, at the optimum, the seller induces a positive-measure set of early-arriving buyers (with intermediate valuations) to wait till period two. Because such buyers contribute to raising the average valuation of the period-two buyers, the seller then has incentives to maintain a higher price in period two, thus reducing the outside option of the early-arriving buyers. In contrast, those early-arriving buyers with either a very low or very high valuation are induced to contract immediately in period one by entering into a long-term (option) contract.

The paper is related to a few recent contributions. Ely, Garrett, and Hinnosaar (2015) consider

a sequential screening problem similar to the one in Deb and Said (2015) but where the seller has limited capacity. They show that overbooking (selling more units than capacity), as practiced by airlines and other service providers, may be an optimal response to the commitment problem. By overbooking, the seller biases the allocation problem against the late-arriving buyers, thus incentivizing early purchases by early-arriving buyers. Garrett (2014) considers a related problem with time-varying valuations, similar to the one in Battaglini (2005) but where agents arrive stochastically over time and have private information about their arrival dates. It shows that the seller needs to provide early-arriving buyers with additional incentives to join immediately, and that all types, even the lowest, must receive a positive information rent. To induce immediate contracting, the seller must tighten the incentives constraints over time, so that a buyer with given valuation, but arriving later, receives less favorable terms. By contrast, Deb (2014) and Garrett (2013) consider related intertemporal problems for durable goods when the seller has full commitment over the entire time horizon. Armstrong and Zhou (2015) show how a seller may use different dynamic (option) contracts to deter buyers from exploring the prices at other sellers, thereby influencing the distribution of outside options and leading to greater rent extraction.

Finally, Deb and Said (2015) also consider a two-sided limited commitment problem whereby a buyer entering into a long-term contract in period one can exit the contract in period two and re-enter the market anonymously as a late-arriving buyer. Optimal contracts in this environment can be significantly different than in the case where re-entry is not feasible. Such differences illustrate well the importance of various commitment and observability assumptions for the predictions the theory delivers about equilibrium contracts and the induced distortions.

The seventh paper in the Symposium by **Miao and Zhang (2015)** considers a continuous-time contracting environment with complete information, but limited commitment, both from the agent and the principal. It proposes a novel duality approach that permits analytical solutions. The paper considers a consumption-insurance problem similar to the one examined in Thomas and Worrall (1988), Kocherlakota (1996), Alvarez and Jermann (2000), Ligon, Thomas, and Worrall (2002), and Ljungqvist and Sargent (2004). These previous works assume a discrete-time setup. The approach typically followed to identify properties of optimal contracts is to use the agent's continuation utility as a state variable. Optimality is then summarized by a Hamilton–Jacobi–Bellman (HJB) equation with state constraints. However, such an equation typically does not admit analytical solutions and

may be difficult to analyze even numerically, due to the fact that the space of continuation utility is endogenous. The contribution of Miao and Zhang (2015) is in showing that sharp analytical predictions can be obtained by considering the dual problem and casting the analysis in continuous time. The dual has the advantage of being an unconstrained problem, and delivers a simple linear HJB equation subject to free-boundary conditions.

The paper provides a dynamic programming characterization of the dual problem using individual income as the usual state variable and the cumulative of the Lagrange multipliers associated with the intertemporal participation constraints as co-state variables. When the principal can commit to a dynamic risk-sharing contract but the agent can walk away at any period (in the literature, this case is referred to as "one-sided limited commitment"), the only co-state variable is the one associated with the agent's participation constraints. In particular, when the principal and the agent have the same discount factor, the co-state variable is also equal to the ratio of the two players' marginal utilities. The agent's current income and the ratio of marginal utilities constitute the state variables of the dual problem. When the principal's and the agent's discount rates differ, the co-state variable must be adjusted by the difference in the discount rates, but the adjusted co-state variables continue to be current income and the ratio of marginal utilities.

When neither the agent nor the principal can commit (in the literature, this case is referred to as "two-sided limited commitment"), the analysis parallels the one in the one-sided limited commitment case, but with two co-state variables, corresponding to each of the two players' participation constraints.

The paper illustrates well the advantages of the dual approach by considering two examples where neither autarky nor full risk sharing are outcomes of optimal contracts. The first example is a continuous-time version of the discrete-time models analyzed in Thomas and Worrall (1988), Krüger and Uhlig (2006), and Chapter 19 of Ljungqvist and Sargent (2004). In this example, the principal is a risk-neutral planner and the agent is a household with a constant-relative-risk-aversion utility function. Only the agent has limited commitment: at any instant, the agent can "walk away" from the contract, in which case autarky prevails in the continuation. While in the discrete-time model of Ljungqvist and Sargent (2004), the agent and the principal have the same discount factor and income follows an independently and identically distributed process with a finite state space, in the Miao and Zhang (2015) example, the agent and the principal are allowed to have different

discount factors and income follows a geometric Brownian motion. The corresponding dynamics are quite different in the two models. In particular, while the agent is eventually fully insured in Ljungqvist and Sargent (2004), in Miao and Zhang (2015), the agent is never fully insured. Such differences are not merely a consequence of continuous time. They originate in the fact that the income process is unbounded and continuous in Miao and Zhang (2015), whereas it is bounded and discrete in Ljungqvist and Sargent (2004).

The second example features an environment similar to the one in the first example, but where the principal may also renege on the contract and guarantee herself a constant outside option equal to zero. This problem is a continuous-time version of the problems analyzed in Kocherlakota (1996), Alvarez and Jermann (2000), Ligon, Thomas, and Worrall (2002), and Chapter 20 of Ljungqvist and Sargent (2004). In these previous works, depending on parameters' values, the optimal contract features full risk sharing, autarky (no risk sharing), or limited risk sharing. In particular, autarky is the only sustainable allocation when the discount factor is sufficiently small, whereas full risk sharing can be achieved when the discount factor is sufficiently large. By contrast, in Miao and Zhang (2015), only limited risk sharing can be sustained in equilibrium, irrespective of parameters' values. Once again, the difference is not a mere consequence of continuous time. It also reflects the fact that the income process is unbounded in Miao and Zhang (2015). That full risk sharing can never be an equilibrium outcome then follows from the fact that, for extreme income realizations, the autarky value exceeds any constant-utility level from full risk sharing (the same result holds in discrete-time, when the income process is unbounded). That autarky can never be an equilibrium outcome is perhaps more interesting and appears to be more specific to the continuous-time setting. In fact, in the discrete time approximation of the model, autarky is the only equilibrium outcome if the nonstationary income process is not too volatile and/or the principal and the agent are sufficiently impatient.

The paper also offers interesting comparative statics with respect to the agent's relative risk aversion, the volatility of the income process, and the discount factors. While not all the results are equally surprising, the method used to arrive to the results is elegant and applicable to other settings. That the method permits sharp analytical predictions is quite impressive given that models with heterogenous discount rates and/or two-sided limited commitment are notoriously intractable.

Related to Miao and Zhang (2015) is the eighth paper in the Symposium by **Williams (2015)**.

This paper studies a tractable continuous-time principal-agent model with both moral hazard and hidden savings. The framework is an extension of Holmstrom and Milgrom (1987). As in Williams (2011), the results are established using the *stochastic maximum principle* and showing that the first-order approach to the agent’s incentives problem is valid in the framework under examination. The key contribution relative to earlier continuous-time works is in the explicit characterization of the optimal policies (effort and consumption). This is made possible by the combination of the assumptions that (a) the production technology and asset accumulation are linear in the agent’s effort and in the shocks, (b) both the principal and agent have exponential preferences over consumption, and (c) the agent has quadratic financial costs of effort.

The paper offers various useful insights about the dynamics of the “labor wedge” (the labor/leisure margin) and the “intertemporal wedge” (the consumption/savings margin). As noticed in previous work, when the principal can control the agent’s consumption, both the labor/leisure margin and the consumption/savings margin are distorted. With hidden savings, instead, the consumption/savings margin is undistorted, and, as a result, the labor wedge is larger. Finally, the paper shows how the optimal contracts can be implemented by providing the agent with a constant share of output (or equity share), a constant flow payment, and a constant tax on savings. Contrary to Holmstrom and Milgrom (1987), under the optimal contract, total compensation is not linear in the firm’s cash flows, but it remains linear in the logarithm of the agent’s promised utility.

The paper provides a complete and explicit treatment of the first-best contract, the contract for hidden effort but with no hidden savings, and the contract for hidden effort and hidden savings. Such a complete characterization favors interesting comparisons. That incentive compatibility under hidden savings can be verified analytically is also a major plus relative to previous work.

The ninth paper in the Symposium by **Strulovici and Szydlowski (2015)** considers a family of single-agent dynamic problems where uncertainty is generated by a Brownian motion. The authors consider both stochastic control and optimal stopping problems, such as those analyzed in the agency models of Bergemann and Strack (2015a) and Kruse and Strack (2015) in this Symposium. In the stochastic control setting, the agent’s objective is to control a one-dimensional, time-homogeneous diffusion process in order to maximize the expected rewards. The approach taken in the paper is to analyze the Hamilton–Jacobi–Bellman (HJB) equation associated to the problem under examination. The paper establishes a verification theorem that guarantees that the

value function of the control problem satisfies this equation. These and other related stochastic optimization problems frequently arise in economics, in applications that range from optimal growth models to multi-armed bandit problems. After heuristically deriving a HJB equation, one would like to rely on some general result guaranteeing that this equation possess a unique solution that coincides with the value function. Unfortunately, general existence results for HJB equations are very limited.

The authors make progress by imposing (i) Lipschitz continuity and linear growth conditions with respect to the state variable that hold uniformly across the controls, (ii) continuity with respect to the control, and (iii) a uniform condition of non-vanishing volatility. Under these conditions, Strulovici and Szydlowski (2015) establish that the value function is the unique smooth solution to the HJB equation. In turn, they show that the HJB equation can be used to express a candidate optimal strategy in terms of the value function and its derivatives. Finally, Strulovici and Szydlowski (2015) provide conditions under which the candidate solution indeed leads to an admissible optimal control.

While the smoothness of value functions has been extensively studied in the optimization literature, the contribution here is in the predictions that are made possible by restricting attention to a *one-dimensional* state. In these problems, the value function can be characterized solely by using *ordinary* differential equations rather than *partial* differential equations. In particular, it is enough for the volatility to be of locally bounded variation. Another advantage of the structure considered in the paper is that it permits one to identify sufficient conditions for the existence of admissible optimal controls (in most of previous works, existence has to be checked on a case-by-case basis).

The last two papers in the Symposium consider environments in which, in principle, the designer could run a static mechanism, but where a dynamic mechanism that sequentially discloses information to the agents can improve the relevant objective (the social efficiency of the voting system in the first paper and the revenue to the seller in the second paper).

In particular, **Bognar, Börgers, and Meyer-ter-Vehn (2015)** consider a canonical voting setting in which a finite number of agents vote over binary alternatives. Each voter has a privately-known valuation for the alternatives and all voters share a common prior. The authors are interested in finding a (possibly sequential) voting mechanism that maximizes expected social welfare.

Voting is costly and hence the mechanism must maximize expected utilities, net of voting costs.

While values are private and independent, payoffs are not restricted to be quasilinear. As a result, the designer cannot use transfers to provide incentives to the agents. This assumption implies that, in contrast to the quasilinear case where the first best can be implemented under fairly general conditions (see, e.g., Bergemann and Välimäki (2010)), in the present environment, the possibility of implementing the first-best outcome cannot be taken for granted. The paper shows that the first best can be implemented under a uniform common prior (i.e., when all alternatives are ex-ante equally likely). However, when the prior is sufficiently asymmetric, one can construct examples where the first best cannot be implemented.

The optimal mechanism is naturally one of sequential voting. Because voting is costly, the planner wants to incur the cost of consulting additional voters only when posterior beliefs remain sufficiently close to the prior (note that sequential mechanisms are optimal even in the absence of incentive constraints). The restriction to binary alternatives in turn implies that abstentions can be counted in favor of any arbitrarily chosen default alternative. Sequential voting terminates as soon as one of the alternatives obtains a supermajority of the votes. Importantly, the supermajority threshold declines over time as the probability that the remaining uncast votes reverse the leading alternative declines.

The analysis is closely related to the earlier work of Gershkov and Szentes (2009). The present model differs from Gershkov and Szentes (2009) in a number of important dimensions. First, the environment is one of pure private values, whereas Gershkov and Szentes (2009) consider a Condorcet Jury Theorem environment with pure common values. Second, in the present paper, voters are already informed and the cost is with respect to the act of voting rather than with respect to information acquisition.

In an extension, the paper considers costly voting with common – rather than private – values and rediscovers the impossibility of first-best implementation, which is familiar from the information acquisition literature. This result mirrors a similar distinction between private and interdependent values in static mechanism design with information acquisition, as observed in Bergemann and Välimäki (2002).

A central question in dynamic mechanism design is how to tailor incentives to the sequential arrival of private information. In most works, the process governing the gradual arrival of information is assumed to be exogenous. However, in many environments of interest, the agents' (private)

information is at least in part controlled by the designer. Indeed, in many auction environments, the seller explicitly controls how much information the bidders have about the characteristics of the object on sale. In the federal offshore wildcat oil tract auctions, for example, prior to bidding, interested firms are permitted to gather information about the lease value and their drilling costs using seismic information, but no on-site drilling is allowed (see Porter (1995)). Similarly, Genesove (1993) reports that in wholesale used car auctions, different auctioneers adopt remarkably different rules as to how potential bidders may inspect a used car prior to bidding.

The last paper in the Symposium by **Bergemann and Wambach (2015)** considers an auction setting with a single indivisible good and independent private values. The designer's objective is to maximize revenues by designing a sequential mechanism that jointly determines the allocation of the good, the payments, and the information disclosed to the bidders.

The novelty relative to previous works (most notably, Esó and Szentes (2007)) is in allowing for disclosure policies that release information *sequentially* over time. The paper shows that the optimal mechanism can be implemented through an ascending-price auction in which all the losing bidders learn perfectly their true valuations whereas the winning bidder receives only limited information about his valuation that is censored from below. In other words, the winning bidder only learns that his true value is sufficiently high to win against the losing bidders. Equivalently, at each point in time, a bidder is either informed that the price has reached the bidder's true value, in which case the bidder drops out, or he receives no information (which amounts to learning that the true valuation is strictly higher than the current price). The disclosure process (and the auction itself) ends when there is only one bidder left who has not yet learned his true value for the object. Losing bidders pay nothing, whereas the winning bidder pays a price that is linked to the expectation of the winning bidder's true valuation conditional on it being above the stop-out price (the precise formula depends on whether or not the bidders possess private information prior to participating in the auction).

The analysis in Bergemann and Wambach (2015) builds on previous work by Bergemann and Pesendorfer (2007) and Esó and Szentes (2007). Bergemann and Pesendorfer (2007) consider the optimal design of an information structure in a single-unit auction with independent private values. In contrast to the present work, Bergemann and Pesendorfer (2007) (a) restrict attention to static disclosure policies whereby each bidder receives a single signal prior to bidding and (b) assume

that bidders do not possess private information prior to entering the auction. The environment in Bergemann and Pesendorfer (2007) can thus be viewed as an instance of Bayesian persuasion, as formulated in Kamenica and Gentzkow (2011), but one in which (i) there are many receivers (the bidders) and (ii) the sender (the seller) designs not only the disclosure rule but also the game among the receivers. The present work follows Eső and Szentes (2007) in that it allows the bidders to possess private information prior to participating in the mechanism. In fact, the ascending auction of Bergemann and Wambach (2015) implements the same final allocation of the optimal handicap auction of Eső and Szentes (2007). However, because the ascending auction reveals less information than the handicap auction to the winning bidder, the allocation is implemented under more stringent participation constraints. In particular, in Eső and Szentes (2007), the participation constraint of each bidder is satisfied at the interim level—i.e., prior to participating in the mechanism—but not ex post. In fact, in the handicap auction, all bidders make a payment prior to entering the auction (this payment can be thought of as the price for the handicap). As a result, all losing bidders’ ex-post participation constraints are violated. In contrast, in Bergemann and Wambach (2015), losing bidders pay nothing (and hence their ex-post participation constraints are trivially satisfied) whereas the winning bidder pays (in expectation) more than in the handicap auction, but less than his ex-post expected valuation. Crucially, the larger payment by the winning bidder is sustained by disclosing him less information than in the handicap auction. In other words, the sequential auction of Bergemann and Wambach (2015) implements the same allocation rule as the handicap auction of Eső and Szentes (2007), but under a more stringent solution concept, posterior equilibrium (see Green and Laffont (1987) for a definition).

In both Bergemann and Wambach (2015) and Eső and Szentes (2007), the information disclosed to the bidders is a function of the bidders’ reported initial types and of *incremental information* that is orthogonal to the bidders’ true initial types. While it is always without loss of generality to describe the bidders’ true valuations as functions of their initial types and of orthogonal incremental information (as implied by the Integral Transform Probability theorem; see also Pavan, Segal, and Toikka (2014) for a discussion of how this applies to arbitrary processes), restricting the disclosure policy to depend on the reported initial types as opposed to the true ones can be with loss of generality. For example, Li and Shi (2015) show that *static* disclosure policies that depend on the true initial types and on orthogonalized shocks (equivalently, that depend on the true ex-post

valuations) can improve the seller's revenue relative to the handicap auction (and thus also relative to the ascending auction analyzed in Bergemann and Wambach (2015)). Extending the analysis of optimal sequential disclosure rules by allowing such rules to depend on the ex-post valuations represents an important future step for this line of research.

4 Avenues for future research

The extensive scope of the contributions in the Symposium illustrates well the progress that the dynamic analysis of contracts and mechanisms has made since the early contributions by Baron and Besanko (1984) and Courty and Li (2000). As the interest in this research program continues to grow, we expect significant further progress in the years ahead. We conclude this Introduction with a brief discussion of some of the questions that remain open and of the challenges that the field is facing.

Commitment, Dynamic Arrivals and Departures, and Participation Constraints. Most of the literature assumes that principals and agents arrive contemporaneously and that principals can commit to exclude those agents who declined earlier offers from trading in future periods. As shown in Skreta (2015) in this Symposium, relaxing such commitment assumptions may uncover novel aspects of optimal mechanisms, especially when the agents' private information changes over time. Furthermore, as discussed in Deb and Said (2015), in many environments of interest, the population of agents changes over time, as new agents arrive and existing ones depart from the marketplace. For example, Deb and Pai (2013) and Mierendorff (2014) derive optimal mechanisms when the private information of the buyers is multi-dimensional, namely their valuation and the deadline by which they must buy or consume the good to be sold. The additional considerations and incentive constraints that enter through the deadline typically require a dynamic version of ironing and the subsequent implementation of the mechanism leads to "biased" auctions.

Importantly, arrivals and departures may not be observable by the principal, implying that those agents arriving earlier can effectively choose between the menus meant for them and those meant for agents expected to arrive and exit at different dates. In other words, agents can choose *when to* enter in and/or exit from a mechanism. This choice can be particularly valuable to the agents when

their private information is expected to change over time, as shown first in Garrett (2013). Such dynamic considerations in turn may lead to larger information rents and to a different structure of the intertemporal distortions (relative to those under observable arrivals and departures). Garrett (2013) characterizes the optimal contracts for the provision of a durable good in a setting in which agents arrive stochastically *and* their valuations change over time. Garrett (2014) in turn extends the analysis to contracts for the provision of non-durable services, as in Battaglini (2005). Similarly, Bergemann and Strack (2015b) offer an analysis of stationary contracts in the continuous-time environment of Bergemann and Strack (2015a); the requirement of stationarity is meant to restrict the principal to offering identical contracts to all future incoming generations. A common element in these very recent contributions is the requirement that the participation constraints must be satisfied at all periods. Notably, these works restrict attention to settings with either a single agent or multiple agents but without strategic interactions. Extending the analysis of mechanism design to dynamic settings in which (a) multiple agents arrive stochastically over time and compete in each period, and (b) the principal lacks the commitment to treat agents arriving at different dates differently is expected to generate interesting novel effects.

Dynamic Disclosure and Information Management. In much of the literature on dynamic mechanism design the focus is on the sequential elicitation of the private information that the agents receive over time. There are two closely related issues that arise with the sequential elicitation of information. First, how much of the information gathered by the principal should be revealed to other participating agents. In the context of auctions with resale, Calzolari and Pavan (2006a) show that partial disclosure is necessary for revenue maximization. In the context of sequential voting, Gershkov and Szentes (2009) show that it may be optimal to reveal neither the outcome nor the number of votes cast. In the context of repeated first-price auctions, Bergemann and Hörner (2014) show that disclosing past bids may favor implicit collusion, thus lowering the seller's revenue. Second, as discussed in Bergemann and Wambach (2015) in the present Symposium, when the primitive information of each agent is endogenous, an optimal mechanism must specify the speed of information disclosure. In recent work, Li and Shi (2015) analyze how, in a static environment, the principal optimally discloses information so as to balance revenue and efficiency considerations. The framework features a single buyer (or, equivalently, many buyers in the absence

of capacity constraints). As mentioned above, the key difference relative to Esó and Szentes (2007) and Bergemann and Wambach (2015) is that Li and Shi (2015) consider a larger class of disclosure policies in which the disclosed information may depend directly on the true state (in contrast, Esó and Szentes (2007) and Bergemann and Wambach (2015) restrict the disclosed information to depend on (a) the agent’s reported initial type and (b) an innovation that is orthogonal to the initial private information). A lot of work remains to be done about the structure of optimal disclosure policies in environments with competing agents.

Beyond the specific settings of sequential screening or sequential auctions, a more general question is how to manage information dynamically over time. The recent work on Bayesian persuasion—to use the language of Kamenica and Gentzkow (2011)—can naturally be extended to dynamic environments. For example, Ely (2015) considers a dynamic model of Bayesian persuasion where the sender controls an entire stochastic process—in the case of Ely (2015), a Poisson process. In related work, Che and Hörner (2015) study the design of a dynamic recommendation system in the presence of social learning about the quality of a new product. The endogeneity of the agents’ private information is also the driving force in the analysis of dynamic taxation with learning by doing of Makris and Pavan (2015) and in the design of dynamic matching with experimentation of Fershtman and Pavan (2015). We expect further progress will come in the years ahead from bridging the dynamic persuasion with the dynamic mechanism design literature.

Beyond Quasi-Linearity and the First-Order Approach. As is typical in the literature on mechanism design, most of the papers in this Symposium consider environments with quasilinear payoffs. There are many settings though in which either monetary transfers are not used (as in certain organizations or in the allocation of certain goods, such as kidneys) or the agents’ preferences are not approximated well by a quasilinear function. Extending the analysis of dynamic mechanisms to non-quasilinear settings is essential to understand many problems of interest, including intertemporal consumption smoothing, taxation, and insurance. This is the focus of the recent dynamic public finance literature as discussed in the Introduction. In related recent contributions, Luz (2015) analyzes dynamic competitive insurance markets and Hörner and Guo (2015) study the design of optimal dynamic mechanisms in a model with a finite Markov chain and without transfers.

A difficulty in settings without quasilinear payoffs is that local incentive constraints and point-

wise optimization conditions typically fail to characterize the optimal contracts. This is an area where significant progress remains to be done that we expect will continue to attract attention in the years ahead.⁶

Empirical Work. Concurrently with the advancement of theoretical work, dynamic mechanisms have been adopted recently in various real-life environments to facilitate economic transactions. Most airlines, for example, now use dynamic pricing algorithms. Furthermore, the recent widespread diffusion of digital and decentralized data has favored novel applications such as dynamic pricing by hotels, restaurants, and car services such as Uber and Lyft. We therefore expect that new theoretical questions will continue to arise in conjunction with the development of new real-life, real-time, economic applications.

A better understanding of the theoretical properties of dynamic mechanisms should also allow for a more productive interaction between theoretical and empirical work. Just as static contract theory helped gain insights into the empirical regularities of various markets (see, e.g., Chiappori and Salanié (2003) and Einav, Finkelstein, and Levin (2010) for excellent surveys of the empirical literature), we expect the recent developments in the dynamic mechanism design literature to help empirical work in the analysis of industries. Among important early contributions here are Chiappori, Salanié, and Valentin (1999) on the dynamics of wage formation under learning, and Hendel and Lizzeri (2003) on the role of commitment in dynamic contracts with data from life insurances. More recent contributions pertain to mobile phone services, health clubs, sequential auctions and insurance markets, see for example Miravete (2003) and Grubb and Osborne (2015) for service plans, Aron-Dine, Cullen, Einav, and Finkelstein (2015) and Handel, Hendel, and Whinston (2015) for health insurance markets, or Einav, Farronato, Levin, and Sundaresan (2013) for internet auctions.

⁶Recent contributions aimed at extending the analysis to settings in which the first-order approach is not necessarily valid include Battaglini and Lamba (2015) and Garrett, Pavan, and Toikka (2015).

References

- ACEMOGLU, D., M. GOLOSOV, AND A. TSYVINSKI (2011): “Political Economy of Ramsey Taxation,” *Journal of Public Economics*, 95, 467–475.
- AKAN, M., B. ATA, AND J. DANA (2015): “Revenue Management by Sequential Screening,” *Journal of Economic Theory*, forthcoming.
- ALBANESI, S., AND C. SLEET (2006): “Dynamic Optimal Taxation with Private Information,” *Review of Economic Studies*, 73, 1–30.
- ALVAREZ, F., AND U. JERMANN (2000): “Efficiency, Equilibrium, and Asset Pricing with Risk of Default,” *Econometrica*, 68, 775–797.
- ARMSTRONG, M., AND J. ZHOU (2015): “Search Deterrence,” *Review of Economic Studies*, forthcoming.
- ARON-DINE, A., M. CULLEN, L. EINAV, AND A. FINKELSTEIN (2015): “Moral Hazard in Health Insurance: Do Dynamic Incentives Matter?,” *Review of Economics and Statistics*, forthcoming.
- ATHEY, S., AND I. SEGAL (2013): “An Efficient Dynamic Mechanism,” *Econometrica*, 81, 2463–2485.
- ATKINSON, A., AND R. LUCAS (1992): “On the Efficient Distribution with Private Information,” *Review of Economic Studies*, 59, 427–454.
- BARON, D., AND D. BESANKO (1984): “Regulation and Information in a Continuing Relationship,” *Information Economics and Policy*, 1, 267–302.
- BARON, D., AND R. MYERSON (1982): “Regulating a Monopolist with Unknown Costs,” *Econometrica*, 50, 911–930.
- BATTAGLINI, M. (2005): “Long-Term Contracting with Markovian Consumers,” *American Economic Review*, 95, 637–658.
- (2007): “Optimality and Renegotiation in Dynamic Contracting,” *Games and Economic Behavior*, 60, 213–246.

- BATTAGLINI, M., AND R. LAMBA (2015): “Optimal Dynamic Contracting: The First-Order Approach and Beyond,” Discussion paper, Cornell University and Penn State University.
- BERGEMANN, D., AND J. HÖRNER (2014): “Should Auctions Be Transparent?,” Discussion Paper 1764R, Cowles Foundation for Research in Economics, Yale University.
- BERGEMANN, D., AND M. PESENDORFER (2007): “Information Structures in Optimal Auctions,” *Journal of Economic Theory*, 137, 580–609.
- BERGEMANN, D., AND M. SAID (2011): “Dynamic Auctions: A Survey,” in *Wiley Encyclopedia of Operations Research and Management Science*, ed. by J. Cochran, no. 1757, pp. 1511–1522. Wiley, New York.
- BERGEMANN, D., AND P. STRACK (2015a): “Dynamic Revenue Maximization: A Continuous Time Approach,” *Journal of Economic Theory*, forthcoming.
- (2015b): “Stationary Dynamic Contracts,” Discussion paper, Yale University and UC Berkeley.
- BERGEMANN, D., AND J. VÄLIMÄKI (2002): “Information Acquisition and Efficient Mechanism Design,” *Econometrica*, 70, 1007–1033.
- (2006): “Information in Mechanism Design,” in *Advances in Economics and Econometrics*, ed. by R. Blundell, W. Newey, and T. Persson, pp. 186–221. Cambridge University Press, Cambridge.
- (2010): “The Dynamic Pivot Mechanism,” *Econometrica*, 78, 771–790.
- BERGEMANN, D., AND A. WAMBACH (2015): “Sequential Information Disclosure in Auctions,” *Journal of Economic Theory*, forthcoming.
- BERGER, A., R. MÜLLER, AND S. NAEEMI (2010): “Path-Monotonicity and Truthful Implementation,” Discussion paper, Maastricht University.
- BESANKO, D. (1985): “Multi-Period Contracts Between Principal and Agent with Adverse Selection,” *Economics Letters*, 17, 33–37.

- BESTER, H., AND R. STRAUZ (2000): “Imperfect Commitment and the Revelation Principle: The Multi-Agent Case,” *Economic Letters*, 69, 165–171.
- (2001): “Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case,” *Econometrica*, 69, 1077–1098.
- BOARD, S. (2007): “Selling Options,” *Journal of Economic Theory*, 136, 324–340.
- (2011): “Relational Contracts and the Value of Loyalty,” *American Economic Review*, 101, 3349–3367.
- BOARD, S., AND A. SKRZYPACZ (2015): “Revenue Management with Forward-Looking Buyers,” *Journal of Political Economy*, forthcoming.
- BOGNAR, K., T. BÖRGERS, AND M. MEYER-TER-VEHN (2015): “An Optimal Voting Procedure When Voting Is Costly,” *Journal of Economic Theory*, forthcoming.
- BOLESLAVSKY, R., AND M. SAID (2013): “Progressive Screening: Long-Term Contracting with a Privately Known Stochastic Process,” *Review of Economic Studies*, 80, 1–34.
- BÖRGERS, T. (2015): *An Introduction to the Theory of Mechanism Design*. Oxford University Press, Oxford.
- BULOW, J. (1982): “Durable Good Monopolists,” *Journal of Political Economy*, 90, 314–332.
- BURGUET, R., AND J. SÁKOVICS (1996): “Reserve Prices without Commitment,” *Games and Economic Behavior*, 15, 149–164.
- CAILLAUD, B., AND C. MEZZETTI (2004): “Equilibrium Reserve Prices in Sequential Ascending Auctions,” *Journal of Economic Theory*, 117, 78–95.
- CALZOLARI, G., AND A. PAVAN (2006a): “Monopoly with Resale,” *RAND Journal of Economics*, 37, 362–375.
- (2006b): “On the Optimality of Privacy in Sequential Contracting,” *Journal of Economic Theory*, 130, 168–204.

- CARBAJAL, J., AND J. ELY (2013): “Mechanism Design without Revenue Equivalence,” *Journal of Economic Theory*, 148, 104–133.
- CHE, Y., AND J. HÖRNER (2015): “Optimal Design for Social Learning,” Discussion paper, Columbia University and Yale University.
- CHIAPPORI, P., AND B. SALANIÉ (2003): “Testing Contract Theory: A Survey of Some Recent Work,” in *Advances in Economics and Econometrics - Theory and Applications, Eight World Congress*, ed. by M. Dewatripont, L. Hansen, and P. Turnovsky, pp. 115–149. Cam.
- CHIAPPORI, P., B. SALANIÉ, AND J. VALENTIN (1999): “Early Starters versus Late Beginners,” *Journal of Political Economy*, 107, 731–760.
- CLARKE, E. (1971): “Multipart Pricing of Public Goods,” *Public Choice*, 8, 19–33.
- COURTY, P., AND H. LI (2000): “Sequential Screening,” *Review of Economic Studies*, 67, 697–717.
- CREMER, J. (1995): “Arm’s Length Relationships,” *Quarterly Journal of Economics*, 110, 275–295.
- DASGUPTA, P., P. HAMMOND, AND E. MASKIN (1979): “The Implementation of Social Choice Rules. Some General Results on Incentive Compatibility,” *Review of Economic Studies*, 66, 185–216.
- DEB, R. (2014): “Intertemporal Price Discrimination with Stochastic Values,” Discussion paper, University of Toronto.
- DEB, R., AND M. PAI (2013): “Ironing in Dynamic Revenue Management: Posted Prices and Biased Auctions,” in *Proceedings of the Symposium on Discrete Algorithms (SODA) 2013*, pp. 620–631.
- DEB, R., AND M. SAID (2015): “Dynamic Screening with Limited Commitment,” *Journal of Economic Theory*, forthcoming.
- EDMANS, A., AND X. GABAIX (2011): “Tractability in Incentive Contracting,” *Review of Financial Studies*, 24, 2865–2863.

- EDMANS, A., X. GABAIX, T. SADZIK, AND Y. SANNIKOV (2012): “Dynamic CEO Compensation,” *Journal of Finance*, 67, 1603–1647.
- EINAV, L., C. FARRONATO, J. LEVIN, AND N. SUNDARESAN (2013): “Sales Mechanisms in Online Markets: What Happened to Internet Auctions,” Discussion paper, Stanford University.
- EINAV, L., A. FINKELSTEIN, AND J. LEVIN (2010): “Beyond Testing: Empirical Models of Insurance Markets,” *Annual Review of Economics*, 2, 311–336.
- ELY, J. (2015): “Beeps,” Discussion paper, Northwestern University.
- ELY, J., D. GARRETT, AND T. HINNOSAAR (2015): “Overbooking,” Discussion paper, Northwestern University and University of Toulouse and Collegio Carlo Alberto.
- ESŐ, P., AND B. SZENTES (2007): “Optimal Information Disclosure in Auctions,” *Review of Economic Studies*, 74, 705–731.
- FARHI, E., AND I. WERNING (2013): “Insurance and Taxation over the Life Cycle,” *Review of Economic Studies*, 80, 596–635.
- FERNANDES, A., AND C. PHELAN (2000): “A Recursive Formulation for Repeated Agency with History Dependence,” *Journal of Economic Theory*, 91, 223–247.
- FERSHTMAN, D., AND A. PAVAN (2015): “Dynamic Matching: Experimentation and Cross Subsidization,” Discussion paper, Northwestern University.
- FREIXAS, X., R. GUESNERIE, AND J. TIROLE (1985): “Planning under Incomplete Information and the Ratchet Effect,” *Review of Economic Studies*, 52, 173–191.
- GARRETT, D. (2013): “Incoming Demand with Private Uncertainty,” Discussion paper, Toulouse School of Economics.
- (2014): “Dynamic Mechanism Design: Dynamic Arrivals and Changing Values,” Discussion paper, Toulouse School of Economics.
- GARRETT, D., AND A. PAVAN (2012): “Managerial Turnover in a Changing World,” *Journal of Political Economy*, 120(879-925).

- (2015): “Dynamic Managerial Compensation: A Variational Approach,” *Journal of Economic Theory*, forthcoming.
- GARRETT, D., A. PAVAN, AND J. TOIKKA (2015): “Robust Predictions in Dynamic Screening,” Discussion paper, Northwestern University and University of Toulouse and MIT.
- GENESOVE, D. (1993): “Adverse Selection in the Wholesale Used Car Market,” *Journal of Political Economy*, 101, 644–665.
- GERSHKOV, A., AND B. MOLDOVANU (2009a): “Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach,” *American Economic Journal: Microeconomics*, 2, 98–168.
- (2009b): “Learning about the Future and Dynamic Efficiency,” *American Economic Review*, 99, 1576–1588.
- (2012): “Optimal Search, Learning and Implementation,” *Journal of Economic Theory*, 147, 881–909.
- (2014): *Dynamic Allocation and Pricing: A Mechanism Design Approach*. MIT Press.
- GERSHKOV, A., B. MOLDOVANU, AND P. STRACK (2014): “Revenue Maximizing Mechanisms with Strategic Customer and Unknown Demand: Name-Your-Own-Price,” Discussion Paper <http://ssrn.com/abstract=2527653>, Hebrew University and University of Bonn and UC Berkeley.
- GERSHKOV, A., AND B. SZENTES (2009): “Optimal Voting Scheme with Costly Information Acquisition,” *Journal of Economic Theory*, 144, 1895–1920.
- GOLOSOV, M., A. TSYVINSKI, AND M. TROSHKIN (2015): “Redistribution and Social Insurance,” Discussion paper, Princeton University and Yale University and Cornell University.
- GREEN, J. (1987): “Making Book Against Oneself,” *Quarterly Journal of Economics*, 98, 785–796.
- GREEN, J., AND J. LAFFONT (1979): *Incentives in Public Decision Making*. North-Holland, Amsterdam.

- (1987): “Posterior Implementability in a Two Person Decision Problem,” *Econometrica*, 55, 69–94.
- GROVES, T. (1973): “Incentives in teams,” *Econometrica*, 41, 617–631.
- GRUBB, M., AND M. OSBORNE (2015): “Cellular Service Demand: Biased Beliefs, Learning and Bill Shock,” *American Economic Review*, forthcoming.
- GUESNERIE, R., AND J. LAFFONT (1984): “A Complete Solution to a Class of Principal-Agent Problems with an Application to the Control of a Self-Managed Firm,” *Journal of Public Economics*, 25, 329–369.
- GUL, F., H. SONNENSCHN, AND R. WILSON (1986): “Foundations of Dynamic Monopoly and the Coase Conjecture,” *Journal of Economic Theory*, 39, 155–190.
- HANDEL, B., I. HENDEL, AND M. WHINSTON (2015): “Equilibria in Health Exchanges: Adverse Selection vs. Reclassification Risk,” *Econometrica*, forthcoming.
- HART, O., AND J. TIROLE (1988): “Contract Renegotiation and Coasian Dynamics,” *Review of Economic Studies*, 55, 509–540.
- HENDEL, I., AND A. LIZZERI (2003): “The Role of Commitment in Dynamic Contracts: Evidence from Life Insurance,” *Quarterly Journal of Economics*, 118, 299–327.
- HINNOSAAR, T. (2015): “Calendar Mechanisms,” Discussion paper, Collegio Carlo Alberto.
- HOLMSTROM, B., AND P. MILGROM (1987): “Aggregation and Linearity in the Provision of Intertemporal Incentives,” *Econometrica*, 55.
- HÖRNER, J., AND Y. GUO (2015): “Dynamic Mechanisms with Money,” Discussion paper, Yale University and Northwestern University.
- KAKADE, S., I. LOBEL, AND H. NAZERZADEH (2013): “Optimal Dynamic Mechanism Design and the Virtual Pivot Mechanism,” *Operations Research*, 61, 837–854.
- KAMENICA, E., AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590–2615.

- KAPICKA, M. (2013): “Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks,” *Review of Economic Studies*, 80, 1027–1054.
- KOCHERLAKOTA, N. (1996): “Implications of Efficient Risk Sharing without Commitment,” *Review of Economic Studies*, 63, 595–609.
- (2005): “Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation,” *Econometrica*, 73, 1587–1621.
- KRÄHMER, D., AND R. STRAUZ (2011): “The Benefits of Sequential Screening,” University of Bonn and Humboldt University of Berlin.
- KRÜGER, D., AND H. UHLIG (2006): “Competitive Risk Sharing Contracts with One-Sided Commitment,” *Journal of Monetary Economics*, 53, 1661–1691.
- KRUSE, T., AND P. STRACK (2015): “Opyimal Stopping with Private Information,” *Journal of Economic Theory*, forthcoming.
- KUMAR, P. (1985): “Consistent Mechanism Design and the Noisy Revelation Principle,” Ph.D. thesis, Stanford University.
- LAFFONT, J., AND J. TIROLE (1986): “Using Cost Observation to Regulate Firms,” *Journal of Political Economy*, 94, 614–641.
- LAFFONT, J., AND J. TIROLE (1988): “The Dynamics of Incentive Contracts,” *Econometrica*, 59, 1735–1754.
- LAFFONT, J., AND J. TIROLE (1990): “Adverse Selection and Renegotiation in Procurement,” *Review of Economic Studies*, 57, 597–625.
- LI, H., AND X. SHI (2015): “Discriminatory Information Disclosure,” Discussion paper, University of British Columbia and University of Toronto.
- LIGON, E., J. THOMAS, AND T. WORRALL (2002): “Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies,” *Review of Economic Studies*, 69, 209–244.

- LIU, Q., K. MIERENDORFF, AND X. SHI (2015): “Auctions with Limited Commitment,” Discussion paper, Columbia University and UCL and University of Toronto.
- LJUNGQVIST, L., AND T. SARGENT (2004): *Recursive Macroeconomic Theory*. MIT Press, Cambridge, Massachusetts, second edition edn.
- LUZ, V. F. (2015): “Dynamic Competitive Insurance,” Discussion paper, University of British Columbia.
- MAKRIS, M., AND A. PAVAN (2015): “Taxation under Learning by Doing: Incentives for Endogenous Types,” Discussion paper, Northwestern University.
- MASKIN, E., AND J. RILEY (1984): “Monopoly with Incomplete Information,” *Rand Journal of Economics*, 15, 171–196.
- MCAFEE, R., AND D. VINCENT (1997): “Sequentially Optimal Auctions,” *Games and Economic Behavior*, 18, 246–276.
- MIAO, J., AND Y. ZHANG (2015): “A Duality Approach to Continuous-Time Contracting Problems with Limited Commitment,” *Journal of Economic Theory*, forthcoming.
- MIERENDORFF, K. (2014): “Optimal Dynamic Mechanism Design with Deadlines,” Discussion paper, UCL.
- MIRAVETE, E. (2003): “Choosing the Wrong Calling Plan? Ignorance and Learning,” *American Economic Review*, 93, 297–310.
- MIRRLEES, J. (1971): “An Exploration in the Theory of Optimum Income Taxation,” *Review of Economic Studies*, 38, 175–208.
- MUSSA, M., AND S. ROSEN (1978): “Monopoly and Product Quality,” *Journal of Economic Theory*, 18, 301–317.
- MYERSON, R. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6, 58–73.
- PAVAN, A., I. SEGAL, AND J. TOIKKA (2014): “Dynamic Mechanism Design: A Myersonian Approach,” *Econometrica*, 82, 601–653.

- PORTER, R. (1995): “The Role of Information in U.S. Offshore Oil and Gas Lease Auctions,” *Econometrica*, 63, 1–27.
- RAHMAN, D. (2011): “Detecting Profitable Deviations,” Discussion paper, University of Minnesota.
- RILEY, J., AND W. SAMUELSON (1981): “Optimal Auctions,” *American Economic Review*, 71, 381–392.
- RIORDAN, M., AND D. SAPPINGTON (1987): “Information, Incentives and Organizational Mode,” *Quarterly Journal of Economics*, 102, 243–264.
- ROCHET, J.-C. (1987): “A Necessary and Sufficient Condition for Rationalizability in a Quasi-Linear Context,” *Journal of Mathematical Economics*, 16, 191–200.
- SAID, M. (2011): “Sequential Auctions with Randomly Arriving Buyers,” *Games and Economic Behavior*, 73, 236–243.
- (2012): “Auctions with Dynamic Populations: Efficiency and Revenue Maximization,” *Journal of Economic Theory*, 147, 2419–2438.
- SANNIKOV, Y. (2013): *Advances in Economics and Econometrics: Tenth World Congress* chap. Contracts: The Theory of Dynamic Principal-Agent Relationships and the Continuous-Time Approach, pp. 89–124. Cambridge University Press.
- SKRETA, V. (2006): “Sequentially Optimal Auctions,” *Review of Economic Studies*, 73, 1085–1111.
- (2015): “Optimal Auction Design under Non-Commitment,” *Journal of Economic Theory*, forthcoming.
- SKRZYPACZ, A., AND J. TOIKKA (2015): “Mechanisms for Repeated Trade,” *American Economic Journal: Microeconomics*, forthcoming.
- STANTCHEVA, S. (2014): “Optimal Income Taxation with Adverse Selection in the Labour Market,” *Review of Economic Studies*, 81, 1296–1329.
- STOKEY, N. (1981): “Rational Expectations and Durable Goods Pricing,” *Bell Journal of Economics*, 12, 112–128.

- STRULOVICI, B. (2010): “Learning While Voting: Determinants of Collective Experimentation,” *Econometrica*, 78, 933–971.
- STRULOVICI, B., AND M. SZYDŁOWSKI (2015): “On the Smoothness of Value Functions and the Existence of Optimal Strategies in Diffusion Models,” *Journal of Economic Theory*, forthcoming.
- THOMAS, J., AND T. WORRALL (1988): “Self-Enforcing Wage Contracts,” *Review of Economic Studies*, 55, 541–554.
- VICKREY, W. (1961): “Counterspeculation, Auctions and Competitive Sealed Tenders,” *Journal of Finance*, 16, 8–37.
- WILLIAMS, N. (2011): “Persistent Private Information,” *Econometrica*, 79, 1233–1274.
- (2015): “A Solvable Continuous Time Dynamic Principal-Agent Model,” *Journal of Economic Theory*, forthcoming.
- WILSON, R. (1993): *Nonlinear Pricing*. Oxford University Press, Oxford.