Dynamic Contracting*

Alessandro Pavan†

Northwestern University and CEPR

August 2025

Abstract

Dynamic contracting plays a central role in many environments—for instance, the sale of goods and services to consumers whose preferences evolve through learning, experimentation, or habit formation; the taxation of workers whose productivity changes with learning-by-doing; the provision of services on platforms with stochastic entry and exit of buyers and sellers; and the matching of agents whose values and attractiveness are gradually revealed through past interactions. This article surveys several strands of the recent dynamic mechanism design literature, distills a few lessons, and points to promising directions for future research.

JEL classification: D82

Keywords: dynamic mechanism design, experimentation, limited commitment, revenue management, financial frictions, robustness.

*For useful comments and suggestions, I thank Laura Doval, George Georgiadis, Harry Pei, Alex Smolin, Bruno Strulovici, and specially Larry Samuelson. I also thank my co-authors on various projects on dynamic mechanism design, Daniel Fershtman, Daniel Garrett, Miltos Makris, Ilya Segal, and Juuso Toikka, with whom I had the pleasure to work over the years, and from whom I learned enormously. Their influence is evident throughout the manuscript. Finally, I thank Rithik Khanna for excellent research assistance. The usual disclaimer applies.

†Department of Economics, 2211 N. Campus Drive. Email: apavan31@gmail.com

1 Introduction

Designing contracts in dynamic settings where information, preferences, trading opportunities, technology, and productivity evolve over time plays an important role in many economic problems. Firms devote significant resources to the writing of employment contracts (especially for top employees) outlining how tasks and compensation will adjust to changes in the business landscape such as the entry of new competitors, variations in regulation, and the arrival of new technologies. Online sellers hire experts to design dynamic pricing algorithms that respond to variations in estimated demand, consumers' purchasing history, and changes in input prices. Lenders include sophisticated covenants in their contracts to structure borrowers' responses to information about their projects' profitability, and consequently their ability to repay the loans. Insurance companies design plans that account for how deductibles and limits on out-of-pocket expenses shape agents' usage of such plans and ultimately the market price of the services the insurers pay for. Platforms design their matching protocols with the consideration that users learn about their preferences for interacting with partners over time, with influence from past matches. Local and federal governments design tax codes accounting for their effect on workers' and firms' incentives to invest in human capital, and how, in turn, changes in labor demand and supply shape the dynamics of workers' productivity through learning-by-doing.

The above are just a few examples of environments where dynamic contracting plays a fundamental role. A key difficulty in designing contracts in such environments stems from the fact not only must incentives be provided to the relevant parties for the revelation of private information over time, but the latter's evolution is often endogenous due to experimentation, learning-by-doing, and habit formation. Furthermore, in many relevant settings, the feasibility of future allocations is shaped by contractual choices made in earlier periods, which makes the design problem inherently non–time-separable.

Dynamic mechanism design offers a toolkit based on research in game theory, contract theory, operations research, and computer science to study these dynamic problems by extending static techniques to settings with evolving private information, strategic interactions, and possibly non time-separable payoffs. It permits one to address questions such as the following:

- How do market power distortions evolve under profit-maximizing contracts? In particular, when do they diminish over time and disappear in the long run?
- How do service providers respond to population dynamics, by which buyers, sellers, and goods enter and exit the market stochastically at different points in time?
- How does limited commitment shape the structure and efficiency of long-term relationships?
- To what extent do the features of optimal mechanisms persist under limited knowledge of the environment under non-Bayesian approaches?

Answering these questions entails addressing complex trade-offs—for example, between targeting efficiency and rent extraction, or between providing incentives to privately informed players and ensuring contract credibility under limited commitment. Doing so requires sophisticated tools, such as recursive formulations of the agents' incentive problem and the identification of inter-temporal conditions that guarantee contract credibility. These tools not only facilitate answering the types of questions mentioned above but also provide concrete insights for both positive and normative analysis in many real-world situations.

This article surveys the theory of dynamic mechanism design, revisiting earlier contributions, and then focusing on a few recent lines of research.[1] The exposition is neither technical nor exhaustive. It aims to provide a reader new to this literature with an exposure to the key trade-offs that arise when (a) information must be solicited at multiple points in time, (b) limited commitment and financial frictions constrain the relevant allocations, (c) the stochastic arrival of trading opportunities shape the dynamics of trade, and (d) uncertainty about the environment constrains the way contracts respond to contingencies. The selection of the topics covered is subjective and does not give full justice to the literature. For example, the exposition does not cover important topics including the following:

- *dynamic moral hazard* (see, e.g., Sannikov, (2013) for an extension of previous work to continuous time, and Ely et al. (2025) and the references therein for some of the recent developments in this literature);
- *incomplete contracts* (see Hart (2017) for an overview and Che and Sakovics (2021) for some recent developments, focusing on dynamic effects);
- *variations to the durable-good monopolist problem and the Coase conjecture* not covered in Section 5 (see, e.g., Fuchs and Skrzypacz (2010), Board and Pycia (2014), and Nava and Schiraldi (2019) for a few examples);
- *behavioral contract theory* (see the survey by Kőszegi (2014), and Jakobsen (2020) for a discussion of recent contributions);
- *empirical analysis of dynamic contracts* (see, e.g., Hendel (2017) for the role of commitment in dynamic insurance, and the discussion in Ghili et al. (2024) for more recent research on related topics).

Organization. Section 2 contains a simple stylized model that illustrates the intricacies and key trade-offs in optimal mechanism design when information must be solicited at different points in time. Section 3 mentions a few applications. The exposition in these sections is concise (and in many ways incomplete) to make room, in Sections 4-9, for a discussion of a few directions of recent research: population dynamics, limited commitment, disclosure, endogenous types (stemming from experimentation, information acquisition, and learning-by-doing), ex-post participation constraints, financial frictions, and robustness. Section 10 concludes with a view towards future work.

2 Incentive compatibility, virtual surplus, and distortion dynamics

In this section, I introduce a stylized screening model with time-varying private information that abstracts from many of the effects mentioned above (the dependence of the feasibility of future decisions on

[1] See also the Introduction that Dirk Bergemann and I wrote for the Journal of Economic Theory Symposium Issue on Dynamic Contracts and Mechanism Design (Bergemann and Pavan, (2015)), the article I prepared for the Econometric Society 11th World Congress (Pavan, (2017)), and Bergemann and Välimäki (2019)'s survey on dynamic mechanism design.

earlier ones, population dynamics, limited commitment, disclosure, endogenous types, financial frictions, and uncertainty about the environment) but illustrates how incentive compatibility in dynamic settings is shaped by the interaction between the dynamics of the allocations and the connection between earlier and future types. See Pavan et al. (2014) for a general treatment of dynamic incentive compatibility in richer environments.

2.1 Environment and contractual problem

The principal (she) is a seller, the agent (he) is a buyer. Their relationship lasts for T periods, where T can be finite or infinite. Time is discrete and indexed by $t = 1, 2, ..., T$. Both the buyer and the seller have time-additively separable preferences given, respectively, by

$$U^P = \sum_t \delta^{t-1}(p_t - C(q_t)) \quad \text{and} \quad U^A = \sum_t \delta^{t-1}(\theta_t q_t - p_t)$$

where $q_t \in \mathcal{Q} \subset \mathbb{R}$ denotes the quantity exchanged in period t, $\theta_t \in \Theta_t$ denotes the buyer's period-t marginal value for the seller's product, p_t denotes the *total* payment from the buyer to the seller in period t, $\delta \geq 0$ denotes the common discount factor, and $C(q_t)$ denotes the cost to the seller of providing quantity q_t. The function C is strictly increasing, convex, and differentiable.

Let $F \equiv (F_t)$ denote the collection of kernels describing the evolution of the buyer's private information: F_1 denotes the prior distribution over Θ_1 and, for all $t \geq 2$, $F_t(\cdot \mid \theta_{t-1})$ denotes the cumulative distribution function of θ_t given θ_{t-1}.

The events occur as follows:

- At $t = 0$, (prior to entering any negotiations with the principal), the buyer privately learns θ_1.
- At $t = 1$, the seller offers a mechanism $\varphi = (\mathcal{M}, \phi)$ and commits to it. The mechanism consists of a collection of mappings

  $$\phi_t : \mathcal{M}_1 \times \cdots \times \mathcal{M}_t \to \mathcal{Q} \times \mathbb{R},$$

  one for each period. Each of these mappings specifies a quantity-price pair for each possible history of messages $m^t \equiv (m_1, ..., m_t) \in \mathcal{M}_1 \times \cdots \times \mathcal{M}_t$, with $\mathcal{M} \equiv (\mathcal{M})_{t=1}^T$ and $\phi \equiv (\phi_t)_{t=1}^T$. If the buyer refuses to participate in φ, the game ends and both players receive a payoff of zero. If the buyer chooses to participate in φ, he sends a message $m_1 \in \mathcal{M}_1$, receives the quantity $q_1(m_1)$, pays a transfer $p_1(m_1)$, and the game moves to period 2.
- $\cdots$
- At the beginning of each period $t \geq 2$, the buyer privately learns θ_t. He then sends a new message $m_t \in \mathcal{M}_t$, receives the quantity $q_t(m^t)$, pays $p_t(m^t)$ to the principal, and the game moves to period $t + 1$.
- $\cdots$
- At $t = T + 1$ the game is over (when T is finite).

The agent has "deep pockets," meaning he can borrow to pay upfront for the surplus expected in future periods. In turn this last assumption implies that the principal can guarantee that, even if the agent can leave the mechanism after joining in period 1, he will never find it optimal to do so (see the discussion in Section 8).

The principal seeks to design a mechanism and a strategy for the agent that maximize the principal's payoff while being consistent with the agent's rationality. Because the principal can commit, by the Revelation Principle[2], without losing optimality, one can restrict attention to direct mechanisms in which $\mathcal{M}_t = \Theta_t$ for all t and such that the agent finds it optimal to always report truthfully. For simplicity, hereafter, I drop the message spaces and identify such a mechanism directly with the policies $\chi = \langle \mathbf{q}, \mathbf{p} \rangle$ that the direct mechanism induces: For any $t \geq 1$, $q_t : \Theta^t \to \mathcal{Q}$ is the period-t *output policy* and $p_t : \Theta^t \to \mathbb{R}$ is the period-t *payment policy*, with $\Theta^t \equiv \Theta_1 \times \cdots \times \Theta_t$.

2.2 Dynamic incentive compatibility

Given the kernels F, one can think of each period-t type θ_t as originating from the combination of an initial type θ_1 and shocks $\varepsilon \equiv (\varepsilon_s)_{s=2}^T$ drawn independently from a uniform distribution over $(0,1)$. Let $Z \equiv (Z_t)$ denote the collection of the functions describing the realizations of the type process, with $\theta_t = Z_t(\theta_1, \varepsilon)$ for all t. One can then define the *impulse response* of θ_s to θ_1 at history θ^s as

$$I_{(1),s}(\theta^s) \equiv \mathbb{E}\left[\frac{\partial Z_s(\theta_1, \tilde{\varepsilon})}{\partial \theta_1} \mid (Z_\tau(\theta_1, \tilde{\varepsilon}))_{\tau=1}^s = \theta^s\right], \tag{1}$$

where the expectation is over the shocks that, given θ_1, lead to the history θ^s.[3] For example, when θ_t follows an autoregressive process of order 1, i.e., for all t, $\theta_t = \gamma\theta_{t-1} + \varepsilon_t$, with $\gamma \in \mathbb{R}_+$, the impulse response of θ_s to θ_1 at history θ^s is simply $I_{(1),s}(\theta_1, \varepsilon) = \gamma^{s-1}$.

This representation often permits one to identify necessary and sufficient conditions for incentive compatibility in dynamic settings. Let $I_{(t),s}(\theta^s)$ denote the generalization of the impulse response function in (1) to an arbitrary pair of periods (t,s) with $s \geq t$; that is, $I_{(t),s}(\theta^s)$ is the impulse response of θ_s to θ_t at history θ^s. For all t, $\theta^t \in \Theta^t$, and $\hat{\theta}_t \in \Theta_t$, then let

$$D_t(\theta^t; \hat{\theta}_t) \equiv \mathbb{E}\left[\sum_{s \geq t} \delta^{s-t} I_{(t),s}(\tilde{\theta}^s) q_s(\tilde{\theta}^s_{-t}, \hat{\theta}_t) \mid \theta^t\right] \tag{2}$$

denote the expected discounted sum of future output, starting from history θ^t, when the period-t report is $\hat{\theta}_t$ and all future types are reported truthfully, with the discounting combining the time discount factor with the impulse response of future types to the period-t one, and with $\theta^s_{-t} \equiv (\theta^{t-1}, \theta_{t+1}, ..., \theta_s)$ denoting the period-s type history, after excluding the period-t component θ_t. Finally let $V_t^A(\theta^t)$ denote the agent's expected payoff under truth-telling in all periods, evaluated at the type history θ^t. The following result characterizes dynamic incentive compatibility:

[2]See, for example, Myerson (1981).

[3]Tildes are used to denote random variables, or random vectors, over which expectations are computed.

Theorem. *Suppose the process F is regular.*[4] *The mechanism $\chi = \langle \mathbf{q}, \mathbf{p} \rangle$ is incentive compatible if, and only if, (a) for all $t \geq 1$, all θ^{t-1}, $V_t^A(\theta^t)$ is equi-Lipschitz continuous in θ_t with*

$$\frac{\partial V_t^A(\theta^t)}{\partial \theta_t} = D_t(\theta^t; \theta_t) \ a.e. \ \theta_t \in \Theta_t, \tag{ICFOC}$$

and (b), for all $t \geq 1$, all θ^{t-1}, all $\theta_t, \hat{\theta}_t \in \Theta_t$,

$$\int_{\hat{\theta}_t}^{\theta_t} [D_t((\theta^{t-1}, x); x) - D_t((\theta^{t-1}, x); \hat{\theta}_t)]dx \geq 0. \tag{Int-M}$$

Conditions analogous to (ICFOC) and (Int-M) jointly characterize the entire set of implementable allocations also in settings where payments are either absent, or the agent's payoff is not linear in the payments, as in insurance and optimal taxation. In quasilinear settings, there always exist payments $\mathbf{p}$ that guarantee that the (ICFOC) conditions are satisfied at all histories. In such quasilinear environments, (ICFOC) thus imposes restrictions on the payments $\mathbf{p}$, while the integral monotonicity conditions (Int-M) impose restrictions on the output schedule $\mathbf{q}$. Jointly, (ICFOC) and (Int-M) guarantee that one-stage deviations from truthful reporting are suboptimal. When applied to a well-behaved Markov environment such as the one considered here, the same conditions imply incentive compatibility at all histories.[5]

Condition (ICFOC) is the dynamic analog of the familiar envelope condition from static design requiring that payments be appropriately linked to the allocation rule to guarantee that the net payoffs increase with types at an appropriate rate. Condition (Int-M) in turn generalizes the familiar monotonicity conditions for static settings with supermodular payoffs, unidimensional types, and unidimensional decisions. It requires that the derivative of the agent's payoff with respect to his true type be sufficiently monotone in the reported type. In particular, note that (Int-M) holds in the dynamic environment under examination if the net present value of expected future output, *discounted by impulse responses*

$$\mathbb{E}\left[\sum_{s \geq t} \delta^{s-t} I_{(t),s}(\tilde{\theta}^s) q_s(\tilde{\theta}^s_{-t}, \hat{\theta}_t) \mid \theta^t\right],$$

is *nondecreasing* in the current report $\hat{\theta}_t$. Even if payoffs are supermodular, as in the environment under consideration here, output need not be monotone in each of the reported types (a property referred to in the literature as *strong monotonicity*). It suffices that the sensitivity of the payoffs to the true types be sufficiently monotone, on average, on the reported type, where the average is over states and time, and with the average accounting for the imperfect dependence of future types on the current ones as captured by the impulse response functions. This more permissive notion of monotonicity is key with endogenous types, for a report triggering consumption in one period may also trigger the arrival of information reducing consumption in future periods (see the "bandit auctions" application in Pavan et

[4] The technical definition of regularity can be found in Pavan et al. (2014), but essentially, it means that the kernels F satisfy appropriate differentiability and equi-Lipschitz continuity conditions.

[5] An environment is Markov when the agent's incentives in any period depend only on her current true type and his past reports, but not on his past true types. This property is guaranteed, for example, by the payoffs being time-additively separable and the process corresponding to the kernels F being Markov.

al. (2014) and the matching auctions in Fershtman and Pavan (2017, 2022)).

Also note that the envelope conditions in (ICFOC), when used inductively over time, imply that the quantity schedule $\mathbf{q}$ and the lowest period-1 type's expected payoff $V_1^A(\underline{\theta}_1)$ pin down not only each period-1 type's expected payoff $V_1^A(\theta_1)$ but the entire net present value of the payments in each state θ. This stronger version of payment-equivalence is particularly relevant when the agent's utility function is not linear in the payments, as in most managerial compensation, insurance, and taxation models (see, e.g., Farhi and Werning (2013), Garrett and Pavan (2015), and Makris and Pavan (2021)).

2.3 Relaxed program, dynamic virtual surplus, and the first-order approach

The result in Theorem 2.2 is valid no matter the principal's objective. It is particularly useful, however, when used to characterize properties of profit-maximizing mechanisms. Below, I illustrate the approach followed in most applied papers (often referred to as the "first-order," or Myersonian approach). This approach drops (a) the envelope conditions in ICFOC, for $t > 1$, (b) all the integral-monotonicity conditions in (Int-M) for all t, and (c) all the period-1 participation constraints except the one for the lowest period-1 type to obtain the following *relaxed program:*

$$\mathcal{P}^r : \begin{cases} \max\limits_{\chi=\langle \mathbf{q},\mathbf{p}\rangle} \mathbb{E}\left[\sum_t \delta^{t-1}(p_t(\tilde{\theta}^t) - C(q_t(\tilde{\theta}^t))\right] \\ \text{subject to} \\ V_1^A(\underline{\theta}) \geq 0, \\ \text{ICFOC-(1): } V_1^A(\cdot) \text{ abs. cont. with } \frac{dV_1^A(\theta_1)}{d\theta_1} = D_1(\theta_1;\theta_1) \text{ a.e. } \theta_1. \end{cases}$$

After obtaining the solution to the relaxed program, one must verify that the remaining constraints are also satisfied. Using ICFOC-(1), and integrating by parts, one can rewrite the principal's objective as "*Dynamic Virtual Surplus*":

$$\mathbb{E}\left[\sum_t \delta^{t-1}\left(\left(\tilde{\theta}_t - \frac{1-F_1(\tilde{\theta}_1)}{f_1(\tilde{\theta}_1)} I_{(1),t}(\tilde{\theta}^t)\right) q_t(\tilde{\theta}^t) - C(q_t(\tilde{\theta}^t))\right)\right] - V_1^A(\underline{\theta}). \tag{3}$$

The dynamic virtual surplus function is defined as the sum of the principal's and the agent's gross payoffs, adjusted by penalties that reflect the cost to the principal of leaving surplus (in the form of information rents) to the agent. In dynamic environments, these penalties combine the familiar inverse hazard rate from static mechanism design, $[1 - F_1(\theta_1)]/f_1(\theta_1)$, with impulse response functions $I_{(1),t}(\theta^t)$ that capture the link between the type realized in period t and the initial type in period 1. The inverse hazard rate accounts for the trade-off between leaving rents to types above θ_1 (numerator) and providing efficient allocations to type θ_1 in each period (denominator). The impulse response functions, in turn, measure how a distortion in period-t output at history θ^t affects the rents left to initial types above θ_1: the larger the impulses, the greater the rents.

The representation of the principal's payoff as dynamic virtual surplus holds in any incentive-compatible mechanism, regardless of whether the constraints ignored in the relaxed program bind. The representation is instrumental to characterizing the properties of optimal mechanisms and studying the dynamics of distortions due to profit maximization, as we show next.

2.4 Distortion dynamics

In the context of the simple contracting problem described above, the dynamic virtual surplus function can be maximized history by history. Assume, for example, that the seller's cost C is quadratic and that $\mathcal{Q} = \mathbb{R}_+$. The allocation rule that maximizes dynamic virtual surplus is such that, for any t and θ^t

$$q_t(\theta^t) = \max\left\{\theta_t - \tfrac{1-F_1(\theta_1)}{f_1(\theta_1)} I_{(1),t}(\theta^t); 0\right\}. \tag{4}$$

If the constraints excluded from the relaxed program are satisfied, the optimality conditions in (4) provide useful information about how the principal distorts the provision of quantity over time. In particular, they imply (a) "no-distortion at the top" (that is, output $q_t(\theta^t)$ is at the efficient level θ_t if $\theta_1 = \bar{\theta}$, all t, all θ^t_{-1}), (b), if impulse responses are positive (which is always the case when the process satisfies first-order stochastic dominance), downward distortions for all histories for which $\theta_1 < \bar{\theta}$ and, (c) the dynamics of the impulse responses of future types to the initial ones drive distortion dynamics. The smaller the impulse responses, the smaller the distortions. Furthermore, if impulse responses decline, on average, over time, so do the distortions because smaller impulses imply smaller information rents for the high period-1 types, and therefore a lower benefit to distorting the level of output provided by the principal in future periods.

See Makris and Pavan (2025) for a discussion of how distortion dynamics—or wedges—also depend on intermediate periods' impulse responses (i.e., on $I_{(\tau),s}(\theta^s)$ for arbitrary s, τ with $s > \tau$) in settings in which payoffs are not quasilinear, for instance in dynamic taxation and insurance problems.

The simple design problem presented in this section has been extended in numerous directions to capture richer environments. These include strategic interactions among multiple agents; the endogenous evolution of types as a function of past decisions (e.g., through experimentation, learning-by-doing, or habit formation); non–time-separable payoffs; stochastic arrivals and departures of trading opportunities; financial frictions; and settings where the feasibility of future allocations depends on earlier decisions. Some of these extensions are discussed in the remainder of the article.

3 Applications

Dynamic mechanism design has been applied to a wide range of economic problems. A comprehensive summary of the literature is beyond the scope of this section. Instead, I briefly review a few prominent applications to illustrate the broad methodological portability of the theory.

- **Implementation of efficient allocations.** A core concern in dynamic mechanism design is the implementation of efficient allocations in settings where information arrives gradually and the feasibility of future allocations may depend on past actions. Two prominent contributions are Bergemann and Välimäki (2010) and Athey and Segal (2013). The former develops a dynamic analogue of the Vickrey–Clarke–Groves mechanism (Vickrey, (1961), Clarke (1971), and Groves, (1973)), while the latter introduces a dynamic counterpart to the expected externality mechanism (d'Aspremont and Gérard-Varet (1979)). These papers examine how the fundamental trade-offs among efficiency, incentive compatibility, participation, and budget balance in static settings extend to dynamic environments. See also Liu (2014) for an extension of the Bergemann–Välimäki mechanism to interdependent valuations. A growing literature in economics and computer science builds on these tools to analyze problems such as dynamic public good provision, sponsored search, and dynamic matching.

- **Profit maximization**. A large body of work studies optimal profit-maximizing mechanisms in settings where private information evolves over time. Early contributions include Baron and Besanko (1984) and Riordan and Sappington (1987), with subsequent developments by Courty and Li (2000), Battaglini (2005), Eső and Szentes (2007), Board (2007, 2008), and Kakade et al. (2013). Pavan et al. (2014) unify many of these results in a general dynamic contracting framework that allows for multiple agents, a continuum of types, endogenous type evolution, and arbitrary time horizons. While most of the literature adopts a discrete-time framework, continuous-time models have received increasing attention. Williams (2011) offers a tractable continuous-time model of adverse selection, later extended by Prat and Jovanovic (2014) and Bergemann and Strack (2015) to richer environments. More recently, Bloedel et al. (2025b) identify conditions under which the Williams (2011) result is valid but also qualify that the optimal mechanism may be fundamentally different from the one in Williams (2011) when these conditions are violated.

- **Insurance, taxation, and redistribution.** This strand, often referred to as the "new dynamic public finance" literature, investigates optimal tax and transfer policies when productivity evolves over the lifecycle and is privately observed. Building on Mirrlees (1971), researchers have explored whether income tax progressivity should increase or decrease with age, whether capital income should be taxed, and whether simple tax schedules that ignore income history can still generate large welfare gains (see, e.g., Albanesi and Sleet, (2006), Battaglini and Coate (2008), Farhi and Werning (2013), Kapicka (2013), and Golosov et al. (2016)). Most of these models treat productivity as exogenous. More recent work allows for endogenous skill accumulation and studies how education, human capital investment, and learning-by-doing respond to policy (see, e.g., Stantcheva (2017), Makris and Pavan (2021), and Akcigit et al. (2022)). Related contributions examine optimal insurance in dynamic economies with evolving income, preferences, or endowments (see Farinha Luz (2023), and Bloedel et al. (2025a)).

- **Managerial compensation**. This literature examines the design of optimal dynamic contracts when managerial effort or performance is privately observed. Recent work also accounts for changes in managerial productivity due to technological innovations, industry trends, or regulation (e.g., Edmans and Gabaix (2011), Edmans et al. (2012), Garrett and Pavan (2012, 2015), Carroll and Meng (2016), and Makris and Pavan (2025)). As in the public finance literature, agents are risk-averse, and optimal contracts provide both intra- and inter-temporal insurance. These models shed light on whether pay–performance sensitivity should increase with tenure, how managerial risk exposure varies over the business cycle, and whether contracts that maximize shareholder value lead to socially optimal turnover.

- **Dynamic matching**. Recent work applies dynamic mechanism design to platform-mediated matching environments, where agents' preferences and attractiveness are private and evolve over time. Evolution may be exogenous (e.g., due to market shocks) or endogenous (e.g., as agents learn about their preferences). See Gomes and Pavan (2016, 2023, 2024) for a static treatment and Fershtman and Pavan (2017, 2022) for dynamic extensions. When platform design involves dynamic elicitation of private information under capacity constraints—either at the individual level (as with one-to-one matching) or at the platform level (as in advertising markets)—inefficiencies emerge. These results inform policy debates about how taxation, regulation, or subsidies can mitigate distortions created by market power.

- **Revenue management**. See Section 4.

- **Fire sales under limited commitment**. See Section 5.

4 Population dynamics and revenue management

Many contracting environments feature interesting dynamics originating in (i) evolving private information, (ii) the sequential and stochastic arrival of agents and/or objects to the market, and (iii) limited commitment. The first channel was briefly reviewed (in a stylized way) in Section 2. This section concentrates on the second and reviews a few insights that build on, yet depart in interesting ways from, the optimal-stopping problems studied in the revenue-management and operations-research literatures.

A large body of work characterizes welfare- and profit-maximizing mechanisms when trading opportunities arrive randomly over time and agents possess private information about their preferences. Gershkov and Moldovanu (2014) provide an excellent survey.[6] These papers frame the seller's problem around the trade-off between selling immediately to current customers, potentially forgoing buyers arriving later with higher valuations, and delaying sales in the hope of finding buyers with a higher willingness to pay, at the cost of depressing the utility of those customers already in the system (e.g., because of impatience). This trade-off (absent in the simple model of Section 2) precludes the characterization of the optimal mechanism as the solution to a point-wise maximization of dynamic virtual surplus, even when payoffs themselves are time separable. The optimal mechanism, instead, is typically obtained as the solution to an optimal stopping problem, akin to those studied in operation research. Crucially, the operation-research solution, however, ignores two incentive problems at the heart of the economics literature: persuading arriving agents to (a) participate in the mechanism and (b) reveal their private information. The economics literature shows how to internalize both constraints by modifying the optimal stopping rules.

Most of the contributions to this literature assume the agents' private information is constant over time. Garrett (2016, 2017), instead, allows valuations to evolve privately. When the agents' arrivals are also the agents' private information, the seller can no longer extract the entire surplus of low-valuation types because those buyers may delay entry until their values increase. To forestall such delay, optimal mechanisms often feature physical allocations that deteriorate with calendar time. Alternatively, when quality is fixed, posted prices may need to trend upward. Garrett (2016) illustrates how these forces can lead to price paths that dip in anticipation of sales and rise sharply once a promotion period ends. Garrett (2017), instead, considers the design problem of a seller with unconstrained instruments (full mechanism design) and shows how payoffs at arrival must deteriorate with time to incentivize immediate

[6] See also Board and Skrzypacz (2016) of an illustration of the implications of buyers' forward-looking behavior.

participation. See also Bergemann and Strack (2022) for a related analysis in continuous time.

Akan et al. (2015) study sequential screening when agents privately know when their valuations will change. This environment can be re-cast as a special case of Pavan et al. (2014), but doing so typically breaks the Markov structure that favors tractability. When the timing of learning is itself type-dependent, Akan et al. (2015) show how to navigate the complications and derive clear prescriptions for optimal refund policies, upgrade options, and other dynamic price instruments.[7]

Collectively, these studies demonstrate how incorporating private information about stochastic arrivals and evolving valuations into the framework examined in the traditional revenue-management literature brings concrete new insights that illuminate on the design of efficient and profit-maximizing dynamic, incentive-compatible, selling strategies.

5 Limited commitment

The dynamic mechanism design literature typically assumes the designer can commit to her mechanism, with allocation dynamics originating from either evolving private information (as in the simple model of Section 2) or the stochastic arrival and departure of goods and agents (as in the literature surveyed in Section 4). A related literature on dynamic contracting under limited commitment investigates the dynamics of the relevant allocations when the agents' private information is typically static but the principal cannot commit to future decisions.

A key issue in this literature is identifying canonical mechanisms to use to solve for the optimal ones. With commitment, the Revelation Principle provides the answer: The class of direct and incentive-compatible mechanisms is canonical in that any allocation sustainable by the principal offering mechanisms from any given class can also be sustained by having the principal offering incentive-compatible direct mechanisms.[8]

However, with limited commitment, this class is not canonical. A principal learning the agent's initial type may have incentives to renege on her promises and offer a new continuation contract that removes some of the inefficiencies in the earlier contract. Consider the simple framework in Section 2. Recall that the principal distorts output at history $\theta^t \equiv (\theta_1, ..., \theta_t)$ to limit the information rents of period-1 types $\mathring{\theta}_1 > \theta_1$ (these types are more likely than type θ_1 to have a period-t type exceeding θ_t and hence to obtain the rent the principal must leave to type θ_t in period t to induce information revelation in that period). Once the principal learns the agent's period-1 type, the value of these inefficiencies may decline. In particular, when types are perfectly persistent, the principal finds it optimal to supply first-best output at any $t > 1$ after learning the agent's period-1 type. Anticipating this possibility, the agent may strictly prefer to lie in period 1, a phenomenon referred to as "ratcheting" in the contracting literature.

Thus, in general, inducing the agent to report his type truthfully induces a loss of optimality in dynamic settings where the principal lacks commitment. One may conjecture that treating the principal

[7] See also Boleslavsky and Said (2013) for an analysis of progressive screening when agents possess private information about some parameters of the type process over and above the early realizations.

[8] The result has been established for solution concepts such as dominance and Bayes-Nash equilibrium. See Sugaya and Wolitzky (2021) for an extension to dynamic solution concepts, and Epstein and Peters (1999), Peters (2015), Yamashita (2010), and Attar et al. (2025) for games with competing designers.

as a regular player whose behavior must also be incentivized would solve this problem. Namely, one should be able to restrict attention to mechanisms (run by a mediator) that solicit information from the agent(s) privately and send recommendations (about the allocations to implement) to the principal which are followed on path. However, this approach is not useful in the problem under consideration because one must specify the payoffs the principal can attain by disregarding the recommendations, bringing the problem back to its starting point.

Recent years have seen renewed interest in contracting under limited commitment. For example, Doval and Skreta (2022) show that, when the principal can use a mediator to garble the agent's messages and, at any point in time, the agent retains superior information vis-a-vis the principal, then the following is a canonical class of mechanisms: (a) The agent reports his type at the beginning of each period, and (b) the mechanism then spits out a signal for the principal equal to a recommended posterior belief for the principal about the agent's type– this signal also indexes the terms of trade in the given period. In equilibrium, the agent reports truthfully and the principal's actual beliefs coincide with the recommended ones. This class of mechanisms favors an information design interpretation whereby the principal, using a mediator as an intermediary, persuades her future selves.

The idea that a principal lacking commitment may need to slow down the speed of learning traces back to the earlier literature on the ratchet effect (e.g., Freixas et al. (1985), Hart and Tirole (1988), Laffont and Tirole (1988, 1990), and Bester and Strausz (2001), among others). This earlier literature, however, assumes no mediation in the interaction between the principal and the agent. The agent mixing over his reports induces the gradual resolution of uncertainty necessary to slow down the principal's learning. Such a mixing can be dispensed with under mediated communication, as shown first in Bester and Strausz (2007) and generalized in Doval and Skreta (2022).

One of the classical applications of dynamic contracting under limited commitment is the sale of an indivisible good by a seller facing a buyer with fixed (and exogenous) private information about her value. Coase (1972) conjectures that when both players are patient enough, prices converge immediately to the maximum of the seller's marginal cost (equivalently, her value) and the buyer's lowest willingness to pay, so the seller fails to extract any surplus from the buyer. Gul et al. (1986) formalize this problem as a proper game and prove the conjecture. Ausubel and Deneckere (1989) show that, when the buyer's strategy is not restricted to be stationary, and there is no gap between the lowest buyer's valuation and the seller's cost, a folk-theorem type of result holds: As the time interval between successive periods vanishes, the set of seller's equilibrium payoffs spans the entire interval from zero to the static monopoly profit. Recently, Groseclose (2025) qualifies that the validity of the Coase conjecture in the Gul et al. (1986)'s analysis relies not only on the restriction to stationary strategies but also on the assumption that both players discount the future at the same rate. Liu (2022) also revisits the Coase conjecture by considering a more general setting in which the buyer and the seller may possess additional information about each other's beliefs and shows how the division of surplus depends on second-order uncertainty.

Another way in which the conjecture may not be robust is that it is established by restricting the seller to posting a price in each period. While the assumption is inconsequential without a mediator, the validity of extending the result to mediated communication is the subject of recent work. Doval and Skreta (2024a) use the approach in their earlier paper (Doval and Skreta (2022)) to prove that, when the seller can offer only one-period mechanisms, she cannot do better than posting a price.[9] The seller's inability to commit to future terms, however, does not imply that she cannot benefit from offering a dynamic mechanism. For example, a mechanism that solicits information from the buyer, encrypts it, and reveals it to the seller at future periods may incentivize the seller to continue deploying the mechanism even after observing no trade since offering a new mechanism would require granting the buyer additional rents to elicit information again. A recent paper by Brzustowski et al. (2023) shows that indeed dynamic mechanisms of this sort (referred to as "smart contracts" in that paper) permit the seller to appropriate part of the surplus.

Two questions remain unsettled: (1) what is the maximal surplus that the seller can extract from the buyer, and (2) what class of mechanism is fully canonical for this type of problems. A recent paper by Lomys and Yamashita (2025) tackles the first question by partly bypassing the second. Instead of identifying a canonical class of mechanisms and then optimizing the seller's payoff over such a class, the paper investigates properties of sequences of contractual terms (probability of trade and payments) consistent with both players' sequential rationality. As in Brzustowski et al. (2023), the mediator can recommend dynamic mechanisms to the seller. The principal can replace the mediator at any point, reflecting limited commitment. The paper looks at a relaxation of the primitive problem in which the seller's option value (following any rejection) is given by an arbitrary function of the players' information structure at the time the mediator is replaced. The authors show that a certain function can be used to set up a dynamic mechanism design problem whose solution may yield the principal's largest attainable payoff. Follow-up work is still required, but the paper has the potential to finally resolve question (1) above.[10]

To the best of my knowledge, the literature still lacks an answer to question (2). It is unclear what class of mechanisms is fully canonical and what payoffs are attainable under limited commitment. Lomys and Yamashita (2025) work directly with sequences of primitive contractual variables, but while this idea is appealing, it remains unclear what constraints such sequences must satisfy in general. Doval and Skreta (2022)'s idea of turning the problem into a constrained information design one is also appealing. However, for their approach to yield a fully canonical class of mechanisms, one must allow for dynamic mechanisms and asymmetric disclosures. The latter are no longer first-order beliefs, and it remains unclear how to deal with richer disclosures in full generality, especially within a dynamic mechanism that garbles information at multiple periods. These questions are not mere theoretical curiosities, but rather foundational issues at the heart of dynamic mechanism design and information design.

The literature mentioned above looks at environments in which the principal can dispense with any contract offered in previous periods without the agent's consent. An alternative form of limited commitment arises when the players can write long-term contracts, and guarantee their enforcement if one of them wants so. In this case, both players are required to renegotiate, as in the literature pioneered by Dewatripont (1989). Whereas the earlier literature assumes constant types, more recent literature

[9] See also Doval and Skreta (2024b) for an extension to a setting where the buyer's type evolves according to a Markov chain.

[10] The paper also shows that, contrary to Coase's conjecture, the seller can appropriate positive surplus by offering a low initial price followed by a higher price until an endogenously set deadline. The mechanism is different from the one in Brzustowski et al. (2023) but shares with the latter the idea that encrypting the buyer's information and disclosing it to the seller at a future date may discipline the seller, in turn permitting her to appropriate surplus.

permits time-varying private information (see, e.g., Battaglini (2007)). Additionally, earlier papers in the renegotiation literature assume a finite horizon. Extending the analysis to infinite horizons causes complications because one cannot restrict the principal to offering a direct revelation mechanism in the last period and then proceeding by backward induction. Maestri (2017) and Gerardi and Maestri (2020) allow for infinitely many periods, retaining the same renegotiation protocol as in the earlier literature, under which the principal can offer a new mechanism at any point in time. If the agent accepts it, the mechanism replaces the original one and is implemented till both parties accept to replace it with a new one. Strulovici (2017) considers an alternative protocol whereby renegotiation breaks down at any time with an exogenous probability and shows how this protocol may lead to outcomes resembling those under the Coase conjecture.

Most of this literature assumes a single principal contracting with a single agent. Dynamic contracting with limited commitment and multiple agents has been examined primarily in the context of auctions; see, for example, Skreta (2015) and Liu et al. (2019). Ekmekci et al. (2024) study how the ratcheting effects can be overcome in the presence of a continuum of agents. The idea is that deviations to new contracts are punished by the agents playing a different continuation equilibrium in future periods.

Deb and Said (2015) consider a particular form of limited commitment with multiple agents wherein the seller can commit to the dynamic contract she offers to each agent but cannot commit to the contracts she will offer to agents joining the mechanism in future periods. The seller may then optimally induce some of the high-type agents arriving early to postpone their participation in the mechanism to future dates to induce herself to maintain high prices when contracting with consumers arriving late.

To the best of my knowledge, dynamic contracting with multiple principals under limited commitment has not been considered yet and represents an important next step for this literature.

6 Disclosure of exogenous and endogenous information

Closely related to the dynamic mechanism design literature is the growing body of work on information design (also referred to as Bayesian persuasion) investigating optimal information disclosures in various settings. Some recent developments are surveyed in Hörner and Skrzypacz (2017), Bergemann and Morris (2019), and Kamenica (2019). Most of this literature is static. For dynamic extensions, see, for example, Ely (2017), Doval and Ely (2020), and Ely and Szydlowski (2020).

The canonical persuasion model assumes the sender (the designer) has access to any mapping of the state into a distribution over a rich set of signal realizations. The sender seeks to influence the receiver's behavior by shaping his beliefs. However, in many economically relevant environments, the sender must acquire information before passing it on to the receiver—for example by screening another agent. Calzolari and Pavan (2006a, 2006b) were the first to explore this integration of screening and persuasion. In their model, a principal screens the private information of one or more agents before transmitting a garbled version of that information to a third party. More recently, Dworczak (2020) considers the problem of a sender constrained to announcing whether the agent's type is above or below a cutoff. Such mechanisms, in general, need not be optimal. However, they provide useful insights about the interaction between screening and persuasion with aftermarkets.

Bergemann et al. (2025) study a setting where persuasion precedes screening: a seller endowed with information beneficial to a buyer discloses the information to the buyer and then screens the buyer's endogenous willingness to pay before finalizing the terms of trade. With standard interim participation constraints, the principal can extract the full surplus by asking the agent to sign the contract before receiving any private information. By contrast, with ex-post participation constraints, the design problem yields interesting predictions about optimal menu's cardinality—which is significantly smaller than that of the agent's endogenous types in many cases.

A complementary literature focuses on the design of optimal disclosure rules in auctions. Key contributions include Eső and Szentes (2007), Bergemann and Wambach (2015), Nikandrova and Pancs (2017), Li and Shi (2017), and Krähmer (2020). Some of these studies examine optimal information structures within specific auction formats (e.g., Bergemann et al. (2022)), while others characterize optimal mechanisms more broadly (e.g., Krähmer (2020)). A central question in this literature is whether the principal must leave rents to the agents for information disclosed privately to them after the initial contractual stage. Eső and Szentes (2007) argue this is not necessary. However, Pavan et al. (2014) note this conclusion hinges on the validity of the first-order approach. Krähmer and Strausz (2015a) revisit the same model with a finite type space (Eső and Szentes (2007) assume a continuum) and show that the principal may indeed need to pay rents for the agents' future private information, even when the latter is orthogonal to that possessed at the contractual stage. The broader lesson is that when agents acquire information over time, and the principal can control this process, full disclosure is generally suboptimal, as nicely illustrated in Li and Shi (2017). Optimal policies must balance the value of tailoring decisions to more precise information about the relevant state against the additional rents induced by greater transparency.

Deb et al. (2025) contribute to this literature by studying a setting where a privately informed agent communicates with a principal who then seeks to persuade a third party. Unlike Calzolari and Pavan (2006a, 2006b), the third party can observe the agent's message. The principal randomizes over the mapping from messages to allocations and privately informs the agent of the selected lottery's realization before the agent sends his message. Remarkably, this structure replicates the outcomes achievable under private communication (between the agent and the principal) and noisy disclosure (to the third party): the observability constraint is non-binding, as a carefully designed lottery can fully substitute for opacity.

Despite these advances, most of the existing literature features limited dynamics, typically a single round of persuasion followed by a single round of screening, or vice versa. In contrast, many applications involve richer inter-temporal structures. Extending the analysis to fully dynamic settings will allow for a deeper understanding of how information is optimally distributed across stages, how its anticipation affects strategic behavior, and how distortions evolve. These questions remain largely open and present promising directions for future research.

7 Endogenous types: Experimentation, information acquisition, and learning-by-doing

Most of the dynamic mechanism design literature treats the evolution of the agents' private information as exogenous. While this assumption is appropriate for many applications, there are important settings where the type process's endogeneity is central. Consider the sale of experience goods. The buyers' valuations evolve as the result of their consumption, and thus the mechanism itself shapes the agents' future information. Similarly, endogenous type processes are natural in environments with habit formation, addiction, or learning-by-doing.

A few recent contributions explore this dimension. Prominent examples include the efficient sequencing model of Bergemann and Välimäki (2010), the procurement model of Krähmer and Strausz (2011), the sponsored search model of Kakade et al. (2013), the bandit-auction model of Pavan et al. (2014), the dynamic matching model of Fershtman and Pavan (2017, 2022), and the taxation model with learning-by-doing of Makris and Pavan (2021).

The framework developed by Pavan et al. (2014) accommodates decision-dependent processes. The authors apply their model to the sale of experience goods and show that the profit-maximizing mechanism takes the form of a sequence of bandit auctions, where the good is allocated to the bidder with the highest "virtual" Gittins index in each period. This index mirrors that of the classic multi-armed bandit literature but is adjusted downward by an handicap accounting for the principal's cost of leaving informational rents to the agents to incentivize truthful information revelation.

Kakade et al. (2013), and Fershtman and Pavan (2017, 2022) offer alternative approaches to dynamic experimentation with private information. Kakade et al. (2013) study a dynamic allocation problem similar to that of Pavan et al. (2014) but assume that payoffs are separable—either additively or multiplicatively—in the agents' initial types. This separability enables the implementation of the virtual index rule using a mechanism akin to the dynamic pivot mechanism of Bergemann and Välimäki (2010). A key feature of their design is that agents re-report their initial type in every period. This "re-reporting" Markovizes the environment by permitting those agents who lied in the past to correct possible inefficiencies in the allocation of the good due to the initial reports. Separability, in turn, implies that each agent's continuation payoff is either proportional to, or an affine transformation of, her marginal contribution to virtual surplus. This property, along with the fact that the handicaps are determined solely by the initial reports, implies that truthful reporting is incentive-compatible both on and off the equilibrium path, provided agents expect others to report truthfully.[11]

Fershtman and Pavan (2017, 2022) adapt this approach to a dynamic matching environment where a platform intermediates interactions between two sides of a market. Each side is populated by agents whose preferences for potential partners are private and evolve over time. The baseline analysis assumes an exogenous type evolution, whereas the supplement explores the implications of type endogeneity. The model features a mix of vertical differentiation (some agents value interacting with the other side more than others) and horizontal differentiation (relative preferences over specific partners may vary across agents and time). It also incorporates capacity constraints on the number of matches per period. Building on Gomes and Pavan (2016, 2024)'s work on static matching design, the analysis characterizes the properties of both welfare- and profit-maximizing dynamic matching protocols. See also Jullien et al. (2021) and Gomes and Pavan (2023)). These mechanisms are implemented through auctions, in which agents, upon joining, choose a membership status that influences their treatment in subsequent match-

[11] Specifically, each agent is incentivized to induce a virtually efficient allocation in each continuation game, which can be accomplished by (a) correcting possible period-1 lies and then (b) truthfully reporting the current and future type innovations.

ing, and then, in each period, bid for each potential partner. Each bilateral match is assigned a score combining the agents' mutual bids, membership status, and number of past interactions. In each period, the matches with the highest non-negative scores are implemented, subject to capacity constraints. The payment schemes ensure that agents' continuation payoffs reflect their marginal contributions to continuation weighted surplus. Depending on the designer's objective, the latter may coincide with total welfare, the platform's profits, or a combination of the two. The payments are similar in spirit to those of the Generalized Second Price auction used in sponsored search but are adjusted to account for (a) the cost of leaving information rents to the agents and (b) the platform's value of generating information that can be used in later periods. Importantly, although the agents' private values for experimentation differ from the platform's, all agents remain in the mechanism after all histories and truthfully bid their myopic values for all partners at all periods in equilibrium. The analysis is used to shed light on the distortions in matching allocations due to platforms' market power and to derive a few policy implications.

Relatedly, Fershtman and Pavan (2025) examine a dynamic procurement problem where a buyer allocates contracts among a pool of suppliers while recruiting new participants over time. The endogeneity of the suppliers' pool creates a bandit problem with an endogenous and time-varying set of arms. In the profit- and the welfare-maximizing mechanisms, firms truthfully bid their unit cost of supplying the good upon joining and are compensated according to a transfer that accounts for their expected subsequent utilization, with an extra subsidy designed to induce truthful bidding. The analysis reveals that late entrants may earn higher profits than earlier ones, all other things equal. This occurs because, over time, the buyer's beliefs about the remaining potential suppliers become more pessimistic. Consequently, the buyer grants new entrants more failed attempts to deliver a satisfactory service, which contributes, other things equal, to larger profits. See also Fershtman and Pavan (2021) for a bandit problem with an endogenous (time-varying) set of arms applied to recruitment under alternative affirmative action policies.

Makris and Pavan (2021) incorporate learning-by-doing into an otherwise classical dynamic income taxation model with non–quasi–linear payoffs. In their framework, workers' productivity evolves endogenously over the life cycle, generating new trade-offs for optimal tax design. Among other results, the paper shows that learning-by-doing calls for higher wedges—that is, larger distortions relative to the full-information benchmark—and alters the relationship between marginal tax rates and wedges. In a version of the model calibrated to US data, they show that tax reforms which account for learning-by-doing yield substantial welfare gains (in terms of extra consumption over the life cycle). Moreover, a simple, memory-less tax code that conditions tax rates only on current income can approximate the welfare gains of fully optimal codes. To isolate the role of learning-by-doing, the paper compares the optimal code in the baseline calibration to one in an economy identical to the calibrated one except for the productivity process being exogenous, revealing significant differences in both optimal tax rates and redistribution.

Makris and Pavan (2025) offers further theoretical insights. That paper develops a recursive characterization of wedge dynamics in a general class of dynamic economies. The analysis accommodates arbitrarily many periods, an endogenous type process, general planner's preferences for redistribution (captured by non-linear Pareto weights on the agents' lifetime utilities), and rich agents' preferences over consumption (captured by concave utility functions). The resulting wedge formula characterizes distortions' evolution, unifying results in the macro new dynamic public finance literature and the micro dynamic mechanism design literature. The analysis clarifies the forces driving the wedges' evolution and provides a general toolkit for understanding how redistribution, insurance, and information dynamics interact in complex inter-temporal environments.

A different form of endogeneity arises when agents acquire private information before participating in a mechanism (see, among others, Crémer, et al. (1998), Bergemann and Välimäki (2002), Gershkov and Szentes (2009), Tirole (2009), Krähmer and Strausz (2011), and Mensch (2022)), or after a contract offer is on the table (see, among others, Crémer and Khalil (1992), Lewis and Sappington (1997), and Szalay (2009)). In some of these papers, the value of information is only speculative and the optimal mechanism deters its acquisition. In others, information helps identify the efficient allocations, and its acquisition is incentivized. (In this latter case, interesting dynamics arise from the combination of moral hazard and the subsequent screening of acquired information.)

8 Ex-post participation and financial frictions

In the simple model of Section 2, the agent's participation must be incentivized only at the time of contract signing (equivalently, mechanism joining). In some cases, this assumption is not appropriate. In online display advertising markets, publishers may not be able to charge advertisers up-front for the right to display ads to consumers, the profitability of which is typically learned post hoc. This is a classical sequential screening problem with a single allocation but time-varying private information, as studied first by Courty and Li (2000). That paper, however, assumes that participation constraints must be satisfied only ex-ante, i.e., at the time the agents first approach the seller.

Motivated by the prevalence of withdrawal rights, according to which consumers may return a product not meeting their needs, Krähmer and Strausz (2015b) add ex-post participation constraints to the Courty and Li (2000) model. Their paper shows that, if consumers' private information at the contractual stage is only imperfectly correlated with their ex-post willingness to pay, the seller cannot profit from soliciting this information. The optimal contract coincides with the optimal price given the buyers' ex-post information. More generally, the paper shows that no benefits from sequential screening arise whenever differences in ex-ante and ex-post outside options lie below a positive upper bound. Interestingly, endowing consumers with ex-post withdrawal rights limits the seller's ability to extract rents, but need not boost welfare.

The question of whether dynamic contracts do better than static ones, under ex-post participation constraints, is also examined in Bergemann et al. (2020). The paper identifies a necessary and sufficient condition for the optimality of a static contract and shows that, when the buyer's early signal of her ex-post valuation is binary, the optimal dynamic contract may entail a lottery for the agent with the lowest early signal.

Ashlagi et al. (2023) consider a similar problem with multiple goods. The paper shows existence of an optimal mechanism under which each good's allocation occurs when the buyer learns her value for that good. The paper uses this property to show that static mechanisms are suboptimal if the buyer first learns her value for products that are ex-ante less valuable. Under this condition, being able to bundle products is less profitable than the ability to adjust prices dynamically.

Related to the literature on ex-post participation constraints is the literature on dynamic contracting with financial frictions. As mentioned in Section 2, most of the dynamic mechanism design literature assumes that agents have "deep pockets" at the time they are asked to sign a contract with the principal. This assumption permits one to dismiss with periodic and ex-post participation constraints. By asking each agent to post a bond at the contractual stage that would be forfeited if the agent walked away from the relationship, the principal can ensure that only the ex-ante participation constraints bind. This bond, however, can be large, especially for long-lasting relationships and when the agent expects a large surplus in each period.

In many interesting situations, agents are cash-constrained, or unable to borrow at a low rate, so asking them to post large bonds at the contractual stage is infeasible. The structure of the optimal contract then must be modified to account for the difficulty of making agents pay upfront for future information rents and to guarantee that periodic or ex-post participation constraints are satisfied. These requirements imply different distortion dynamics due to market power under the profit-maximizing mechanisms. Consider again the simple model of Section 2. Observe from the formula in (4) that, without financial frictions, the optimal contract entails no distortions from period 2 onwards when types are independent. (To see this, notice that, when types are independent, $I_{1,(t)} = 0$ for all $t > 1$, so output is supplied efficiently after the first period.) This conclusion follows from the fact that, with independent types, future rents necessary to incentivize the revelation of private information received in future periods do not cost anything to the principal who can charge the agents upfront for such rents (uniformly over the agents' period-1 types). This conclusion, however, hinges on "deep pockets" and does not extend to settings with financial frictions. Asking agents to pay upfront for future rents may violate their cash constraints.

Dynamic contracting with liquidity constraints has been examined, among others, by Grillo and Ortner (2018), Fu and Krishna (2019), Krasikov and Lamba (2021), Kräehmer and Strausz (2025), and Liu (2025). The last paper considers a setting in which payments to the agent must be non-negative at all periods but flow payoffs can be negative. Many other papers in this literature, instead, require that flow payoffs themselves be non-negative (as required by periodic participation constraints, which coincide with ex-post participation constraints when evaluated in the last period, as in the model of Krähmer and Strausz (2015b)). Liu (2025) shows that, even if an agent's type (which in the model amounts to the cost of supplying a service) is i.i.d. over time, the optimal mechanism need not treat each period independently, and might instead feature backloading of the payments. Under the optimal contract, the agent can wait for more favorable costs before beginning to work, but when he starts, he is required to continue till the end, lest he loses the right to any payment.

When neither the agents nor the principal have financial constraints, the optimal mechanism is not well defined if the contracting parties have different discount factors and "deep pockets". Accommodating for financial constraints thus also permits the study of the effects of unequal discounting. This is the subject of a recent paper by Krasikov et al. (2023), which examines a canonical dynamic screening problem in which the agent's private information evolves according to a Markov chain and shows that unequal discounting may call for front-loading the agent's rents. Dynamics under the optimal contract are typically cyclical, with infinite memory. Furthermore, global incentive-compatibility constraints may bind under unequal discounting, preventing a characterization of the optimal mechanism as a solution to a relaxed problem.

9 Robustness

Most of the dynamic mechanism design literature models the interaction between the principal and the agents as Bayesian — the principal and the agents share a common prior over the economic environment but differ in their information (the agents' is typically superior). All players maximize subjective expected utility, assign probabilities to unobserved elements of the environment, and revise these probabilities according to Bayes' rule.

A growing body of work has begun to relax these assumptions, exploring the implications of non-Bayesian analysis, limited knowledge of the environment, and robustness to informational uncertainty (see Bergemann and Morris (2012), Carroll (2019), and Brooks and Du (2025) for surveys of different robustness approaches in mechanism design). One influential line of research investigates properties of optimal contracts in environments where the principal lacks precise knowledge of the agents' information (note the difference from assuming the principal knows the distribution from which the agents' information is drawn but not its realization). For instance, Chassang (2013) analyzes a dynamic agency model with moral hazard, adverse selection, and limited liability. He develops a novel approach in which the principal calibrates simple benchmark contracts that would be incentive-compatible without liability constraints and then adapts them to ensure feasibility in the presence of such constraints. The resulting dynamic mechanisms are robust in the sense of being detail-free—i.e., independent of the underlying stochastic process for returns, which need not be i.i.d. or ergodic—and hence deliver robust performance guarantees.

Building on this idea, Chassang and Kapon (2022) study the dynamic allocation of a public resource among liquidity-constrained agents. Transfers are conditioned on successful effort, and agents unable to meet their obligations are ignored. The authors show that limited-liability, detail-free mechanisms that ignore reports from participants who cannot make their promised payments can approximate the performance of the pivot Vickrey–Clarke–Groves mechanism. A combination of cautiousness (to avoid over-reliance on uncertain agents) and forgiveness (to prevent permanent exclusion) yields approximate renegotiation-proofness.

A complementary line of research explores convergence to efficiency in dynamic screening environments under full knowledge of the environment by the principal but limited knowledge by the analyst. Garrett et al. (2025) identify general conditions under which allocations under profit-maximizing contracts converge almost surely to the first-best as the time horizon grows. These conditions are fairly weak and need not coincide with those validating the first-order approach under which convergence to the first-best has been established. Their method bypasses the full characterization of the optimal mechanism by using variational arguments—local perturbations that preserve incentive compatibility—to identify qualitative asymptotic properties of optimal allocations. These perturbations also yield lower bounds on the speed of convergence and permit the identification of certain features of the optimal contracts without fully solving for them. In related work, Battaglini and Lamba (2019) examine dynamic environments with a finite set of types and show that, even when the first-order approach fails, convergence to efficiency holds under approximately optimal contracts. However, a question left open is whether the same conclusion holds if the principal insists on full optimality. Garrett et al. (2025) provide a positive answer provided the type process is an irreducible Markov chain.

Another strand of research investigates price dynamics in settings in which the principal is uncertain about the process governing the agents' evolving private information. Libgober and Mu (2021) consider a monopolist selling an indivisible, durable good to a buyer whose valuation evolves stochastically. The buyer's purchase decision is irreversible, and the seller does not know the arrival process of the buyer's information. The dynamic pricing problem is to maximize the principal's payoff under the worst-case scenario. The main result shows that a constant price maximizes the seller's worst-case payoff over all possible mechanisms and all possible information arrival processes. In a follow-up paper, Li et al. (2024) extend the analysis to a setting with limited seller's commitment. They show that, under broad conditions, equilibrium pricing discourages the buyer from delaying purchase, despite the potential value of waiting for additional information. Malladi (2025) considers a dynamic experimentation setting in which the aggregate demand curve is constant but unknown to the seller. The latter experiments by setting a unit price in each period and observing the value of the demand at that price. The paper characterizes price dynamics yielding the largest profit guarantees when the seller is sequentially rational (in the sense that she acknowledges the implications of current choices on future ones) but non-Bayesian. The strategy yielding the largest guarantee is memory-less and consists in setting a price in each period that maximizes profits in an hypothetical world where the demand curve is the flattest among those running through the price set in the previous period.

Penta (2015) studies belief-robust dynamic mechanism design in the spirit of the robust mechanism design literature (see, e.g., Bergemann and Morris (2012) for an overview). He considers a dynamic environment with many agents, where, despite the presence of a commonly known prior distribution over types, the designer is uncertain about the agents' higher-order beliefs. The paper extends the belief-free implementation concept from static to dynamic settings: A social choice function is robustly partially implementable if it is implementable in perfect Bayesian equilibrium for all admissible belief hierarchies. Penta (2015) shows that this notion of robustness coincides with ex-post incentive compatibility. By contrast, robust full implementation—which requires that all perfect Bayesian equilibria implement the desired social choice function across all belief configurations—is significantly harder to achieve. However, under a relaxed solution concept closer to Bayesian equilibrium, the author derives a recursive characterization of all robustly implementable SCFs. In settings with single-crossing preferences, robust full implementation is achievable if the social choice function satisfies strict ex-post incentive compatibility and a condition limiting payoff interdependencies.

Finally, motivated by applications in computer science (particularly, online advertising), Mirrokni et al. (2020) introduce the concept of non-clairvoyant dynamic mechanisms. The paper studies repeated auctions with time-varying valuations in which neither the designer nor the agents know the process' kernels (i.e., future type distributions). Only the initial period's distribution is common knowledge. A mechanism is non-clairvoyant dynamic incentive compatible if truthful reporting in each period is optimal regardless of an agent's beliefs about future distributions, assuming truthful reporting by others. The authors impose ex-post individual rationality and identify a class of mechanisms that deliver a fraction of the maximum surplus attainable by clairvoyant mechanisms (i.e., those exploiting full knowledge of future type distributions). Their results provide an interesting approach to mechanism design in complex, uncertain dynamic environments.

10 Conclusions

Dynamic mechanism design (dynamic contracting more broadly) has emerged as a powerful framework for positive and normative analysis in economic environments where: (a) information unfolds over time; (b) the feasibility or desirability of future allocations depends on earlier decisions; (c) agents, objects, or trading opportunities arrive stochastically; (d) players face limited commitment; (e) financial frictions preclude payments' front loading; and (f) limited market information calls for robustly optimal mechanisms.

Two main approaches to incentive compatibility in these settings have received particular attention. The first aligns agents' objectives with the planner's via dynamic payments reflecting evolving contributions to total surplus or externalities imposed on others—generalizing the logic of static Vickrey–Clarke–Groves and d'Aspremont–Gérard-Varet mechanisms. The second establishes incentive compatibility by combining dynamic monotonicity conditions on allocation rules with dynamic envelope conditions on payments—building on the Myersonian approach, but accounting for imperfect inter-temporal dependence of private information. Together, these approaches form a versatile toolkit for contract design in a wide range of applications.

Over the past two decades, the literature has deepened our understanding of real-world contracting and inspired innovative interventions. Examples include: showing how sellers can profit from dynamic pricing algorithms that adapt to shifting market conditions (e.g., introducing incentives and private information into operations-research models to price firms' responses to congestion in ride-sharing markets); how vendors can stimulate trade through sophisticated return policies with carefully designed partial refund schemes (e.g., introducing financial frictions and ex-post participation constraints in sequential screening models); why airlines benefit from overbooking followed by auction-like reallocations of scarce seats (e.g., by introducing evolving private information on consumer preferences into models of limited capacity); how platforms leverage user histories to uncover private information critical for matching; how governments can design more efficient and equitable tax systems (e.g., using dynamic models where productivity evolves through learning-by-doing to justify age-dependent income taxes with declining rates); how public and private health insurance can respond to evolving private information on health and income shocks while accounting for the endogenous, time-varying nature of contract usage; and how regulators can curb market power (e.g., by tracking the dynamic evolution of quality/quantity distortions under profit-maximizing contracts).

Despite this progress, several questions remain open. A large share of the literature assumes that private information evolves exogenously. Yet in many contexts of interest, experimentation, learning-by-doing, habit formation, and R&D shape the evolution of types, introducing new forces that may significantly alter conclusions drawn under exogenous-type models. While recent work has begun addressing this gap, the area remains fertile ground for further research.

A parallel line of inquiry concerns limited commitment. Much of this literature abstracts from evolving private information to focus on contract renegotiation dynamics. Extending limited-commitment models to incorporate preference and technology dynamics promises valuable insights for a broader set of applications.

Another common assumption is that principals have complete freedom to design contracts, subject only to incentive and participation constraints. Yet in practice, political economy frictions—such as institutional inertia, lobbying pressures, or legislative rigidities—limit contractual flexibility. Incorporating such frictions can yield more realistic predictions and policy recommendations (see Acemoglu et al. (2011) and Bierbrauer et al. (2021)).

Similarly, much of the literature assumes that monetary transfers are available to provide incentives. Yet in many real-world settings—such as organ transplants, school choice, or the sponsorship of a certain type of academic research—transfers are infeasible. A growing literature explores mechanism design without transfers, in both static and dynamic environments. Notable contributions include Abdulkadiroğlu and Loertscher (2007), Miralles (2012), Kováč et al. (2014), Frankel (2016), Guo and Hörner (2016), and Deb et al. (2018). An especially promising direction is to incorporate hard-evidence acquisition and disclosure, as in Ben-Porath et al. (2019), to improve allocative efficiency.

Finally, the field remains largely theoretical, though empirical interest is rising—particularly in health insurance, consumer retention, and employment relationships (see Ghili et al. (2023)). Most empirical studies still assume stationary private information. Developing tools to identify, estimate, and test models with evolving types would bring empirical work closer to theory, reveal new properties of dynamic contracts, and help shape future research agendas.

References

[1] Abdulkadiroğlu, A., and S. Loertscher (2007): "Dynamic House Allocations," mimeo, Duke University and University of Melbourne.

[2] Acemoglu, D., M. Golosov, and A. Tsyvinski (2011): "Political Economy of Ramsey Taxation," *Journal of Public Economics*, 95 (7-8), 467–475.

[3] Akan, M., B. Ata, and J. D. Dana Jr. (2015): "Revenue Management by Sequential creening," *Journal of Economic Theory* 159 (B), 728–774.

[4] Akcigit, U., D. Hanley, and S. Stancheva (2022): "Optimal Taxation and R&D Policies," *Econometrica*, 90 (2), 645-685.

[5] Albanesi, S., and C. Sleet (2006): "Dynamic Optimal Taxation with Private Information," *Review of Economic Studies*, 73 (1), 1–30.

[6] Ashlagi, I., C. Daskalakis, and N. Haghpanah (2023): "Sequential Mechanisms with Ex Post Individual Rationality," *Operations Research*, 71 (1), 245-258.

[7] Athey, S., and I. Segal (2013): "An Efficient Dynamic Mechanism," *Econometrica*, 81 (6), 2463–2485.

[8] Attar, A., E. Campioni, T. Mariotti, and A. Pavan (2025): "Keeping the Agents in the Dark: Competing Mechanisms, Private Disclosures, and the Revelation Principle," mimeo, University of Toulouse, Universita' degli Studi di Roma Tor Vergata, and Northwestern University.

[9] Ausubel, L. M., and R. J. Deneckere (1989): "Reputation in Bargaining and Durable Goods Monopoly," *Econometrica*, 57 (3), 511–531.

[10] Baron, D. P., and D. Besanko (1984): "Regulation and Information in a Continuing Relationship," *Information Economics and Policy*, 1 (3), 267–302.

[11] Battaglini, M. (2005): "Long-Term Contracting with Markovian Consumers," *American Economic Review*, 95 (3), 637–658.

[12] Battaglini, M. (2007): "Optimality and renegotiation in dynamic contracting," *Games and Economic Behavior*, 60 (2), 213–246.

[13] Battaglini, M., and S. Coate (2008): "Pareto Efficient Income Taxation with Stochastic Abilities," *Journal of Public Economics*, 92 (3-4), 844-868.

[14] Battaglini, M., and R. Lamba (2019): "Optimal Dynamic Contracting: The first-order Approach and Beyond," *Theoretical Economics*, 14 (4), 1435-1482.

[15] Ben-Porath, E., E. Dekel, and B. L. Lipman (2019): "Mechanisms with Evidence: Commitment and Robustness," *Econometrica*, 87 (2), 529-566.

[16] Bergemann, D., F. Castro, and G. Y. Weintraub (2020): "The Scope of Sequential Screening with Ex-post Participation Constraints," *Journal of Economic Theory*, 188, 105055.

[17] Bergemann, D, T. Heumann, and S. Morris (2025): "Screening with Persuasion," Cowles Foundation Discussion Paper No. 2338R3, Yale University.

[18] Bergemann, D., T. Heumann, S. Morris, C. Sorokin, and E. Winter (2022): "Optimal Information Disclosure in Classic Auctions," *American Economic Review: Insights*, 4 (3), 371–388.

[19] Bergemann, D., and S. Morris (2012): "Robust Mechanism Design: The Role of Private Information and Higher Order Beliefs," Clementi, Singapore, World Scientific.

[20] Bergemann, D., and S. Morris (2019): "Information Design: A Unified Perspective," *Journal of Economic Literature*, 57 (1), 44–95.

[21] Bergemann, D., and A. Pavan (2015): "Introduction to Symposium on Dynamic Contracts and Mechanism Design," *Journal of Economic Theory*, 159 (B), 679-701.

[22] Bergemann, D., and P. Strack (2015): "Dynamic Revenue Maximization: A Continuous Time Approach," *Journal of Economic Theory*, 159 (B), 819-853.

[23] Bergemann, D., and P. Strack (2022): "Progressive Participation," *Theoretical Economics*, 17 (3), 1007-1039.

[24] Bergemann, D., and J. Välimäki (2002): "Information Acquisition and Efficient Mechanism Design," *Econometrica*, 70 (3), 1007–1033.

[25] Bergemann, D., and J. Välimäki (2010): "The Dynamic Pivot Mechanism," *Econometrica*, 78 (2), 771–789.

[26] Bergemann, D., and J. Välimäki (2019): "Dynamic Mechanism Design: An Introduction," *Journal of Economic Literature*, 57 (2), 235–274.

[27] Bergemann, D., and A. Wambach (2015): "Sequential Information Disclosure in Auctions," *Journal of Economic Theory*, 159 (B), 1074–1095.

[28] Bester, H., and R. Strausz (2001): "Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case," *Econometrica*, 69 (4), 1077–1098.

[29] Bester, H., and R. Strausz (2007): "Contracting with Imperfect Commitment and Noisy Communication," *Journal of Economic Theory*, 136 (1), 236-259.

[30] Bierbrauer, F. J., P. C. Boyer, and A. Peichl (2021): "Politically Feasible Reforms of Nonlinear Tax Systems," *American Economic Review*, 111 (1), 153-191.

[31] Bloedel, A. W., R. V. Krishna, and O. Leukhina (2025a): "Insurance and Inequality with Persistent Private Information," *Econometrica*, 93 (3), 821-857.

[32] Bloedel, A. W., R. V. Krishna, and B. Strulovici (2025b): "Persistent Private Information Revisited," mimeo, UCLA, Florida State University, and Northwestern University.

[33] Board, S. (2007): "Selling Options," *Journal of Economic Theory*, 136 (1), 324–340.

[34] Board, S. (2008): "Durable-Goods Monopoly with Varying Demand," *Review of Economic Studies*, 75 (2), 391-413.

[35] Board, S. and M. Pycia (2014), "Outside Options and the Failure of the Coase Conjecture," *American Economic Review,* 104 (2), 656-671.

[36] Board, S., and A. Skrzypacz (2016): "Revenue Management with Forward-Looking Buyers," *Journal of Political Economy*, 124 (4), 1046-1087.

[37] Boleslavsky, R., and M. Said (2013): "Progressive Screening: Long-Term Contracting with a Privately Known Stochastic Process," *Review of Economic Studies*, 80 (1), 1–34.

[38] Brooks B. and S. Du (2025): "Informational Robustness in Mechanism Design: A Case for (and Against) the Worst Case," mimeo, University of Chicago and UCSD.

[39] Brzustowski, T., A. Georgiadis-Harris, and B. Szentes (2023): "Smart Contracts and the Coase Conjecture," *American Economic Review*, 113 (5), 1334-1359.

[40] Calzolari, G., and A. Pavan (2006a): "Monopoly with Resale," *Rand Journal of Economics*, 37 (2), 362–375.

[41] Calzolari, G., and A. Pavan (2006b): "On the Optimality of Privacy in Sequential Contracting," *Journal of Economic Theory*, 130 (1), 168–204.

[42] Carroll, G. (2019): "Robustness in Mechanism Design and Contracting," *Annual Review of Economics* 11 (1), 139-166.

[43] Carroll, G., and D. Meng (2016): "Robust contracting with Additive Noise," *Journal of Economic Theory*, 166, 586-604.

[44] Chassang, S. (2013): "Calibrated Incentive Contracts," *Econometrica*, 81 (5), 1935-1971.

[45] Chassang, S., and S. Kapon (2022): "Prior-free Dynamic Allocation under Limited Liability," *Theoretical Economics*, 17 (3), 1109-1143.

[46] Che, Y.-K., and J. Sakovics (2021): "To Trade, or not to Trade, that is the Question: New Roles for Incomplete Contracts in Dynamic Settings," *Journal of the European Economic Association*, 19 (3), 1621–1655.

[47] Clarke, E. H. (1971): "Multipart Pricing of Public Goods," *Public Choice*, 11, 17–33.

[48] Coase, R. H. (1972): "Durability and Monopoly," *The Journal of Law and Economics*, 15 (1), 143-149.

[49] Courty, P., and H. Li (2000): "Sequential Screening," *Review of Economic Studies*, 67 (4), 697–717.

[50] Crémer, J., and F. Khalil (1992): "Gathering Information before Signing a Contract," *American Economic Review*, 82 (3), 566–578.

[51] Crémer, J., F. Khalil, and J.-C. Rochet (1998): "Strategic Information Gathering before a Contract is Offered," *Journal of Economic Theory*, 81 (1), 163–200.

[52] D'Aspremont, C., and L.-A. Gérard-Varet (1979): "Incentives and Incomplete Information," *Journal of Public Economics*, 11 (1), 25-45.

[53] Deb, R., M. M. Pai, and M. Said (2018): "Evaluating Strategic Forecasters," *American Economic Review*, 108 (10), 3057–3103.

[54] Deb, R., M. M. Pai, and M. Said (2025): "Indirect Persuasion," *Journal of Political Economy*, forthcoming.

[55] Deb, R., and M. Said (2015): "Dynamic Screening with Limited Commitment," *Journal of Economic Theory*, 159 (B), 891–928.

[56] Dewatripont, M. (1989): "Renegotiation and Information Revelation Over Time: The Case of Optimal Labor Contracts," *The Quarterly Journal of Economics*, 104 (3), 589-619.

[57] Doval, L., and J. C. Ely (2020): "Sequential Information Design," *Econometrica*, 88 (6), 2575–2608.

[58] Doval, L., and V. Skreta (2022): "Mechanism Design With Limited Commitment," *Econometrica*, 90 (4), 1463-1500.

[59] Doval, L., and V. Skreta (2024a): "Optimal Mechanism for the Sale of a Durable Good," *Theoretical Economics*, 19 (2), 865-915.

[60] Doval, L., and V. Skreta (2024b): "Mechanism Design with Limited Commitment: Markov Environments," *European Economic Review*, 167, 104777.

[61] Dworczak, P. (2020): "Mechanism Design With Aftermarkets: Cutoff Mechanisms," *Econometrica*, 88 (6), 2629-2661.

[62] Edmans, A., and X. Gabaix (2011): "Tractability in Incentive Contracting," *The Review of Financial Studies*, 24 (9), 2865–2894.

[63] Edmans, A., X. Gabaix, T. Sadzik, and Y. Sannikov (2012): "Dynamic CEO Compensation," *The Journal of Finance*, 67 (5), 1603–1647.

[64] Ekmekci, M., L. Maestri, and D.Wei (2024): "Turning the Ratchet: Dynamic Screening with Multiple Agents," mimeo, UC Santa Cruz.

[65] Ely, J. C. (2017): "Beeps," *American Economic Review*, 107 (1), 31-53.

[66] Ely, J. C., G. Georgiadis and L. Rayo (2025): "Feedback Design in Dynamic Moral Hazard," *Econometrica*, 93 (2), 597–621.

[67] Ely, J. C., and M. Szydlowski (2020): "Moving the Goalposts," *Journal of Political Economy*, 128 (2), 468-506.

[68] Epstein, L. G., and M. Peters (1999): "A Revelation Principle for Competing Mechanisms," *Journal of Economic Theory*, 88 (1), 119-160.

[69] Eső, P., and B. Szentes (2007): "Optimal Information Disclosure in Auctions and the Handicap Auction," *Review of Economic Studies*, 74 (3), 705–731.

[70] Farhi, E., and I. Werning (2013): "Insurance and Taxation over the Life Cycle," *Review of Economic Studies*, 80 (2), 596–635.

[71] Farinha Luz, V. (2023): "Optimal Dynamic Insurance Contracts," working paper, University of British Columbia.

[72] Fershtman, D., and A. Pavan (2017): "Pandora's Auctions: Dynamic Matching with Unknown Preferences," *American Economic Review*, 107 (5), 186-190.

[73] Fershtman, D., and A. Pavan (2021): "Soft Affirmative Action and Minority Recruitment," *American Economic Review: Insights*, 3 (1), 1-18.

[74] Fershtman, D., and A. Pavan (2022): "Matching Auctions," *Rand Journal of Economics*, 53 (1), 32-62.

[75] Fershtman, D., and A. Pavan (2025): "Searching for Arms: Experimentation with Endogenous Consideration Sets," mimeo, Tel Aviv University and Northwestern University.

[76] Frankel, A. (2016): "Discounted Quotas," *Journal of Economic Theory*, 166, 396-444.

[77] Freixas, X., R. Guesnerie, and J. Tirole (1985): "Planning under Incomplete Information and the Ratchet Effect," *Review of Economic Studies*, 52 (2), 173–191.

[78] Fu, S., and R. V. Krishna (2019): "Dynamic Financial Contracting with Persistent Private Information," *Rand Journal of Economics*, 50 (2), 418–452.

[79] Fuchs, W., and A. Skrzypacz (2010): "Bargaining with Arrival of New Traders," *American Economic Review*, 100 (3), 802-836.

[80] Garrett, D. F. (2016): "Intertemporal Price Discrimination: Dynamic Arrivals and Changing Values," *American Economic Review*, 106 (11), 3275-3299.

[81] Garrett, D. F. (2017): "Dynamic Mechanism Design: Dynamic Arrivals and Changing Values," *Games and Economic Behavior*, 104, 595-612.

[82] Garrett, D. F., and A. Pavan (2012): "Managerial Turnover in a Changing World," *Journal of Political Economy*, 120 (5), 879–925.

[83] Garrett, D. F., and A. Pavan (2015): "Dynamic Managerial Compensation: A Variational Approach," *Journal of Economic Theory*, 159 (B), 775–818.

[84] Garrett, D., A. Pavan, and J. Toikka (2025): "Robust Predictions in Dynamic Screening," mimeo Northwestern University.

[85] Gerardi, D., and L. Maestri (2020): "Dynamic Contracting with Limited Commitment and the Ratchet Effect," *Theoretical Economics*, 15 (2), 583-623.

[86] Gershkov, A., and B. Moldovanu (2014): "Dynamic Allocation and Pricing: A Mechanism Design Approach," Cambridge, M.A.: MIT Press.

[87] Gershkov, A., and B. Szentes (2009): "Optimal Voting Scheme with Costly Information Acquisition," *Journal of Economic Theory*, 144 (1), 36-68.

[88] Ghili, S., B. Handel, I. Hendel, M. D. Whinston (2023): "Optimal Long-Term Health Insurance Contracts: Characterization, Computation, and Welfare Effects," *Review of Economic Studies*, 91 (2), 1085–1121.

[89] Golosov, M., M. Troshkin, and A. Tsyvinski (2016): "Redistribution and Social Insurance," *American Economic Review*, 106 (2), 359-386.

[90] Gomes, R., and A. Pavan (2016): "Many-to-many Matching and Price Discrimination," *Theoretical Economics*, 11 (3), 1005–1052.

[91] Gomes, R., and A. Pavan (2023): "Two-Sided Markets and Matching Design," in *Online Matching-Based Market Design*. ed. by F. Echenique, N. Immorlica, and V. V. Vazirani. Cambridge, U.K.: Cambridge University Press, 484-508.

[92] Gomes, R., and A. Pavan (2024): "Price Customization and Targeting in Matching Markets," *Rand Journal of Economics*, 55 (2), 230-265.

[93] Grillo, E., and J. Ortner (2018): "Dynamic Contracting with Limited Liability Constraints," working paper, Collegio Carlo Alberto and Boston University.

[94] Groseclose, T. (2025): "The Coase Conjecture When the Monopolist and Customers have Different Discount Rates," *Review of Industrial Organization*, 66, 349–365.

[95] Groves, T. (1973): "Incentives in Teams," *Econometrica*, 41 (4), 617–631.

[96] Gul, F., H. Sonnenschein, and R. Wilson (1986): "Foundations of Dynamic Monopoly and the Coase Conjecture," *Journal of Economic Theory*, 39 (1), 155–190.

[97] Guo, Y., and J. Hörner (2020): "Dynamic Allocation without Money," Working Paper No. 1133, Toulouse School of Economics.

[98] Handel, B., I. Hendel, and M. D. Whinston (2015): "Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk," *Econometrica*, 83 (4), 1261-1313.

[99] Hart, O. (2017): "Incomplete Contracts and Control," *American Economic Review*, 107 (7), 1731–1752.

[100] Hart, O. D., and J. Tirole (1988): "Contract Renegotiation and Coasian Dynamics," *Review of Economic Studies*, 55 (4), 509–540.

[101] Hendel, I. (2017): "Dynamic Selection and Reclassification Risk: Theory and Empirics," in *Advances in Economics and Econometrics: Eleventh World Congress*. ed. by B. Honoré, A. Pakes, M. Piazzesi, and L. Samuelson. Cambridge, U.K.: Cambridge University Press, 99-129.

[102] Hörner, J., and A. Skrzypacz (2017): "Learning, Experimentation, and Information Design," in *Advances in Economics and Econometrics: Eleventh World Congress*. ed. by B. Honoré, A. Pakes, M. Piazzesi, and L. Samuelson. Cambridge, U.K.: Cambridge University Press, 63-98.

[103] Jakobsen, A. M. (2020): "A Model of Complex Contracts," *American Economic Review*, 110 (5), 1243–1273.

[104] Jullien, B., A. Pavan, and M. Rysman (2021): "Two-sided Markets, Pricing, and Network Effects," *Handbook of Industrial Organization*, 4 (1), 485-592.

[105] Kakade, S. M., I. Lobel, and H. Nazerzadeh (2013): "Optimal Dynamic Mechanism Design and the Virtual-Pivot Mechanism," *Operations Research*, 61 (4), 837–854.

[106] Kamenica, E. (2019): "Bayesian Persuasion and Information Design," *Annual Review of Economics*, 11, 249-272.

[107] Kapička, M. (2013): "Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First-Order Approach," *Review of Economic Studies*, 80 (3), 1027-1054.

[108] Kőszegi, B. (2014): "Behavioral Contract Theory," *Journal of Economic Literature*, 52 (4), 1075-1118.

[109] Kováč, E., D. Krähmer, and T. Tatur (2014): "Optimal Stopping in a Principal-agent Model with Hidden Information and no Monetary Transfers," working paper, Universität of Bonn.

[110] Krähmer, D. (2020): "Information Disclosure and Full Surplus Extraction in Mechanism Design," *Journal of Economic Theory*, 187, 105020.

[111] Krähmer, D., and R. Strausz (2011): "Optimal Procurement Contracts with Pre-Project Planning," *Review of Economic Studies*, 78 (3), 1015-1041.

[112] Krähmer, D., and R. Strausz (2015a): "Ex-post Information Rents in Sequential Screening," *Games and Economic Behavior*, 90, 257-273.

[113] Krähmer, D., and R. Strausz (2015b): "Optimal Sales Contracts with Withdrawal Rights," *Review of Economic Studies*, 82 (2), 762-790.

[114] Krähmer, D., and R. Strausz (2025): "Dynamic Screening with Liquidity Constraints," *Economic Theory*, 79 (4), 1421-1453.

[115] Krasikov, I., and R. Lamba (2021): "A Theory of Dynamic Contracting with Financial Constraints," *Journal of Economic Theory*, 193, 105196.

[116] Krasikov, I., R. Lamba, and T. Mettral (2023): "Implications of Unequal Discounting in Dynamic Contracting," *American Economic Journal: Microeconomics*, 15 (1), 638-692.

[117] Laffont, J.-J., and J. Tirole (1988): "The Dynamics of Incentive Contracts," *Econometrica*, 56 (5), 1153–1175.

[118] Laffont, J.-J., and J. Tirole (1990): "Adverse Selection and Renegotiation in Procurement," *Review of Economic Studies*, 57 (4), 597–625.

[119] Lewis, T. R., and D. E. M. Sappington (1997): "Information Management in Incentive Problems," *Journal of Political Economy*, 105 (4), 796-821.

[120] Li, H. and X. Shi (2017): "Discriminatory Information Disclosure," *American Economic Review*, 107 (11), 3363-3385.

[121] Li, Z., J. Libgober, and X. Mu (2024): "Sequentially Optimal Pricing under Informational Robustness," working paper, Columbia University, University of Southern California, and Princeton University.

[122] Libgober, J., and X. Mu (2021): "Informational Robustness in Intertemporal Pricing," *Review of Economic Studies*, 88 (3), 1224-1252.

[123] Liu, H. (2018): "Efficient Dynamic Mechanisms in Environments with Interdependent Valuations: The Role of Contingent Transfers," *Theoretical Economics*, 13 (2), 795-829.

[124] Liu, H. (2022): "Robust Predictions in Coasian Bargaining," *American Economic Review: Insights* 4 (2), 209–22.

[125] Liu, Q., K. Mierendorff, X. Shi, and W. Zhong (2019): "Auctions with Limited Commitment," *American Economic Review*, 109 (3), 876-910.

[126] Liu, Y., (2025): "Dynamic Reward Design," mimeo, Northwestern University.

[127] Lomys, N. and T. Yamashita (2025): "A Mediator Approach to Mechanism Design with Limited Commitment," mimeo, University of Toulouse.

[128] Maestri, L. (2017): "Dynamic Contracting under Adverse Selection and Renegotiation," *Journal of Economic Theory*, 171, 136-173.

[129] Makris, M., and A. Pavan (2021): "Taxation under Learning by Doing," *Journal of Political Economy*, 129 (6), 1878-1944

[130] Makris, M., and A. Pavan (2025): "Wedge Dynamics with Evolving Private Information," working paper, University of Kent and Northwestern University.

[131] Mensch, Jeffrey. 2022. "Screening Inattentive Buyers." *American Economic Review* 112 (6): 1949–84.

[132] Malladi, S. (2025): "Guarantees in Price Experimentation," mimeo Northwestern University, Kellogg School of Management.

[133] Miralles, A. (2012): "Cardinal Bayesian Allocation Mechanisms Without Transfers," *Journal of Economic Theory*, 147 (1), 179–206.

[134] Mirrlees, J., (1971): "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies,* 38, 175–208.

[135] Mirrokni, V., R. Paes Leme, P. Tang, and S. Zuo (2020): "Non-Clairvoyant Dynamic Mechanism Design," *Econometrica*, 88 (5), 1939-1963.

[136] Myerson, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6 (1), 58–73.

[137] Nava, F., and P. Schiraldi (2019): "Differentiated Durable Goods Monopoly: A Robust Coase Conjecture," *American Economic Review,* 109 (5), 1930–1968.

[138] Nikandrova, A., and R. Pancs (2017): "Conjugate Information Disclosure in an Auction with Learning," *Journal of Economic Theory*, 171, 174-212.

[139] Pavan, A. (2017): "Dynamic Mechanism Design: Robustness and Endogenous Types" in *Advances in Economics and Econometrics, Eleventh World Congress*. ed. by B. Honoré, A. Pakes, M. Piazzesi, and L. Samuelson. Cambridge, U.K.: Cambridge University Press, 1-62.

[140] Pavan, A., I. Segal, and J. Toikka (2014): "Dynamic Mechanism Design: A Myersonian Approach," *Econometrica*, 82 (2), 601–653.

[141] Penta, A. (2015): "Robust Dynamic Implementation," *Journal of Economic Theory,* 160, 280-316.

[142] Peters, M. (2015): "Reciprocal Contracting," *Journal of Economic Theory*, 158 (A), 102–126.

[143] Prat, J., and B. Jovanovic (2014): "Dynamic Contracts when the Agent's Quality is Unknown," *Theoretical Economics*, 9 (3), 865-914.

[144] Riordan, M. H., and D. E. M. Sappington (1987): "Information, Incentives, and Organizational Mode," *The Quarterly Journal of Economics*, 102 (2), 243–263.

[145] Sannikov, Y. (2013): "Contracts: The Theory of Dynamic Principal-Agent Relationships and the Continuous-Time Approach," in *Advances in Economics and Econometrics: Tenth World Congress*. ed. by D. Acemoglu, M. Arellano, and E. Dekel. Cambridge, U.K.: Cambridge University Press, 89–124.

[146] Skreta, V. (2015): "Optimal Auction Design under Non-Commitment," *Journal of Economic Theory*, 159 (B), 854–890.

[147] Stantcheva, S. (2017): "Optimal Taxation and Human Capital Policies over the Life Cycle," *Journal of Political Economy,* 125 (6), 1931-1990.

[148] Strulovici, B. (2017): "Contract Negotiation and the Coase Conjecture: A Strategic Foundation for Renegotiation-Proof Contracts," *Econometrica*, 85 (2), 585-616.

[149] Sugaya, T., and A. Wolitzky (2021): "The Revelation Principle in Multistage Games," *Review of Economic Studies*, 88 (3), 1503-1540.

[150] Szalay, D. (2009): "Contracts with Endogenous Information," *Games and Economic Behavior*, 65 (2), 586-625.

[151] Tirole, J. (2009): "Cognition and Incomplete Contracts," *American Economic Review*, 99 (1), 265-294.

[152] Vickrey, W. (1961): "Counter-Speculation, Auctions, and Competitive Sealed Tenders," *The Journal of Finance*, 16 (1), 8–37.

[153] Williams, N. (2011): "Persistent Private Information," *Econometrica*, 79 (4), 1233–1275.

[154] Yamashita, T. (2010): "Mechanism Games With Multiple Principals and Three or More Agents," *Econometrica*, 78 (2), 791–801.