CHAPTER 1

# Dynamic Mechanism Design: Robustness and Endogenous Types

**Alessandro Pavan**

This article was prepared for an invited session at the 2015 World Congress of the Econometric Society. Through a unifying framework, I survey recent developments in the dynamic mechanism design literature and then introduce two new areas that I expect will draw attention in the years to come: robustness and endogenous types.

## 1 INTRODUCTION

Long-term contracting plays an important role in a variety of economic problems including trade, employment, regulation, taxation, and finance. Most long-term relationships take place in a "*changing world*," that is, in an environment that evolves (stochastically) over time. Think, for example, of (a) the provision of private and public goods to agents whose valuations evolve over time, as the result of shocks to their preferences or learning and experimentation, (b) the design of multi-period procurement auctions when firms' costs evolve as the result of past investments, (c) the design of optimal tax codes when workers' productivity evolves over time as the result of changes in technology or because of learning-by-doing, (d) the matching of agents whose values and attractiveness is learned gradually over time through private interactions.

Changes to the environment (either due to exogenous shocks, or to the gradual resolution of uncertainty about constant, but unknown, payoffs) are often anticipated at the time of initial contracting, albeit rarely jointly observed by the parties. By implication, optimal long-term contracts must be flexible to

I thank the organizers of the 2015 World Congress of the Econometric Society, and in particular Larry Samuelson, for the invitation, and my discussant Juuso Välimäki for useful comments and suggestions. I also thank my co-authors on related projects, Daniel Fershtman, Daniel Garrett, Miltos Makris, Ilya Segal, and Juuso Toikka, without whom this work would have not been possible. Finally, I thank Laura Doval, Jonas Mishara-Blomberger, and Bela Szabadi for excellent research assistance.

accommodate such changes, while at the same time provide the parties with incentives to share the information they receive over time.

Understanding the properties of optimal long-term contracts is important both for positive and for normative analysis. It permits one to address questions such as: How does the provision of quantity/quality evolve over time under profit-maximizing contracts? How do the dynamics of the allocations under profit maximization compare to their counterparts under welfare maximization? In particular, when do distortions due to profit maximization decrease over time and vanish in the long run? In what environments does the private observability of the "shocks" (i.e., the changes to the environment subsequent to the signing of the initial contract) play no role? When is the nature of the shocks (i.e., whether they are transitory or permanent) relevant for the dynamics of the decisions under optimal contracts?

The last fifteen years have witnessed significant interest in these questions. Important contributions have been made in extending mechanism design tools to economies in which information evolves over time and a stream of decisions is to be made.[1]

In this article, I first provide a brief overview of the recent dynamic mechanism design literature. I then introduce a simple yet flexible framework that I use in the subsequent sections to review some of the recent contributions. Finally, I discuss two new areas that I expect will attract attention in the near future: robustness and endogenous types.

## 1.1 Brief Review of the Dynamic Mechanism Design Literature

This section builds on a recent overview that I prepared with Dirk Bergemann for the Journal of Economic Theory Symposium Issue on Dynamic Contracts and Mechanism Design (Bergemann and Pavan, 2015).

An important part of the dynamic mechanism design literature studies how to implement efficient allocations in dynamic settings with evolving private information. The pioneering contributions in this area are Bergemann and Välimäki (2010) and Athey and Segal (2013). The first paper constructs a dynamic pivot transfer scheme under which, in each period, all agents receive their expected marginal flow contribution to social welfare. The scheme guarantees that, in each period, all agents are willing to remain in the mechanism and report truthfully their incremental information, regardless of their beliefs

[1] Mechanism design has been used in static settings to examine a variety of problems including: auctions (Myerson, 1981; Riley and Samuelson, 1981; Cremer and McLean, 1988; Maskin and Riley, 1989); nonlinear pricing (Mussa and Rosen, 1978; Wilson, 1993); bargaining (Myerson and Satterthwaite, 1983; Ausubel and Deneckere, 1989, 1993); regulation (Baron and Myerson, 1982; Laffont and Tirole, 1986); taxation (Mirrlees, 1971); political economy (Dasgupta et al., 1979; Acemoglu et al., 2011); public goods provision (Vickrey, 1961; Clarke, 1971; Groves, 1973; Green and Laffont, 1979); organization design (Cremer, 1995), and voting (Gibbard, 1973, 1977; Satterthwaite, 1975). The reader is referred to Börgers (2015) for an excellent overview of the static mechanism design literature.

about other agents' past and current types (but provided they expect others to report truthfully).[2] The scheme can be thought of as the dynamic analog of the various Vickrey–Clarke–Groves (VCG) schemes proposed in static environments. The paper by Athey and Segal (2013), instead, proposes a transfer scheme under which each agent's "incentives payment," at each period, coincides with the variation in the net present value of the expected externality the agent imposes on other agents, with the variation triggered by the agent's own incremental information. The proposed scheme can thus be thought of as the dynamic analog of the type of schemes proposed by d'Aspremont and Gérard-Varet (AGV) for static settings. Relative to the dynamic pivot mechanism of Bergemann and Välimäki (2010), the Athey and Segal (2013) mechanism has the advantage of guaranteeing budget balance in each period. Contrary to Bergemann and Välimäki (2010), however, it need not guarantee that agents have the incentives to stay in the mechanism in each period.[3]

A second body of work investigates properties of profit-maximizing mechanisms in settings with evolving private information. Earlier contributions include Baron and Besanko (1984), Besanko (1985), and Riordan and Sappington (1987). For more recent contributions, see, among others, Courty and Li (2000), Battaglini (2005), Eső and Szentes (2007), Board (2007), and Kakade et al. (2013).

Pavan, Segal, and Toikka (2014) summarize the above contributions and extend them to a general dynamic contracting setting with a continuum of types, multiple agents, and arbitrary time horizon. The model allows for serial correlation of the agents' information and for the dependence of this information on past allocations. The approach to the design of optimal mechanisms in Pavan, Segal, and Toikka (2014) can be thought of as the dynamic analog of the approach pioneered by Myerson (1981) for static settings, and subsequently extended by Guesnerie and Laffont (1984), Maskin and Riley (1984), and Laffont and Tirole (1986), among others. This approach consists in first identifying necessary conditions for incentive compatibility that can be summarized in an envelope formula for the derivative of each agent's equilibrium payoff with respect to the agent's type. This formula in turn permits one to express transfers as a function of the allocation rule and thereby to express the principal's objective as virtual surplus (i.e., total surplus, net of handicaps that control for the cost to the principal of leaving the agents information rents). The second step then consists in maximizing virtual surplus across all possible allocation rules, including those that need not be incentive compatible. The final step consists in verifying that the allocation rule that solves the relaxed program, along with the transfer rule required by the necessary

[2] The formal solution concept capturing the above properties is periodic ex-post equilibrium.

[3] See also Liu (2014) for an extension of the Bergemann and Välimäki (2010) mechanism to a setting with interdependent valuations.

envelope conditions, constitute a fully incentive-compatible and individually-rational mechanism. This last step typically involves "reverse-engineering," i.e., identifying appropriate primitive conditions guaranteeing that the allocation rule that solves the relaxed program satisfies an appropriate monotonicity condition.

The approach in Pavan, Segal, and Toikka (2014) – reviewed in Section 4, below – adapts the above steps to a dynamic environment. The cornerstone is a dynamic envelope theorem that yields a formula for the evolution of each agent's equilibrium payoff and that must be satisfied in any incentive-compatible mechanism. This formula combines the usual direct effect of a change in the agent's current type on the agent's utility (as in static mechanism design problems) with novel effects stemming from the effect that a change in the current type has on the distribution of the agent's future types. These novel effects, which are specific to dynamic problems, are summarized by *impulse response functions* that describe how a change in the current type propagates throughout the entire type process. A second contribution of Pavan, Segal, and Toikka (2014) is to show that, in Markov environments, the aforementioned dynamic envelope formula, combined with an appropriate *integral monotonicity condition* on the allocation rule, provides a complete characterization of incentive compatibility. The integral monotonicity condition is the dynamic analog of the monotonicity conditions identified in static problems with unidimensional private information but multidimensional decisions (see, among others, Rochet, 1987; Carbajal and Ely, 2013; and Berger et al., 2010). This condition requires that the allocations be monotone in the reported types "on average," where the average is both across time and states, and is weighted by the impulse responses of future types to current ones.

As in static settings, the Myersonian (first-order) approach yields an implementable allocation rule only under fairly stringent conditions. An important question for the dynamic mechanism design literature is thus the extent to which the predictions identified under such an approach extend to environments where global incentive-compatibility constraints bind. This topic is addressed in two recent papers, Garrett and Pavan (2015) and, Garrett, Pavan, and Toikka (2016).[4] These papers do not fully solve for the optimal mechanisms. Instead, they use variational arguments to identify certain properties of the optimal contracts. More precisely, they use perturbations of the allocation policies that preserve incentive compatibility to identify robust properties of the dynamics of the allocations under optimal contracts. I review this alternative variational approach in Section 5, below.[5]

[4] See also Battaglini and Lamba (2015).

[5] The notion of robustness considered in these papers is with respect to the details of the type process. Robustness with respect to the agents' higher-order beliefs is the topic of a by now rich literature well summarized in the monograph by Bergemann and Morris (2012). The type of problems examined in this literature are typically static. For some recent developments to

Another body of the literature studies the design of efficient and profit-maximizing mechanisms in dynamic settings where the agents' private information is static, but where agents or objects arrive stochastically over time. A recent monograph by Gershkov and Moldovanu (2014) summarizes the developments of this literature (see also Bergemann and Said, 2011; Board and Skrzypacz, 2016; Gershkov et al., 2014; and Said, 2011, 2012). Most of the papers in this literature assume that the agents' information is stationary. Instead, Garrett (2016a, 2016b), Hinnosaar (2016), and Ely et al. (2016) combine dynamics originating from stochastic arrivals with dynamics generated by evolving private information. A recent new addition to this literature is Akan et al. (2015); the paper studies a sequential screening environment à la Courty and Li (2000), but in which different agents learn their valuations at different times, with the timing of learning correlated with the agents' initial valuations.[6]

Dynamic mechanism design has also been applied to study optimal insurance, taxation, and redistribution in the so-called "New Dynamic Public Finance" literature. For earlier contributions, see Green (1987), Atkenson and Lucas (1992), and Fernandes and Phelan (2000). For more recent contributions, see Kocherlakota (2005), Albanesi and Sleet (2006), Farhi and Werning (2013), Kapicka (2013a), Stantcheva (2014) and Golosov et al. (2016).

In all the papers above, the evolution of the agents' private information is exogenous. In contrast, the evolution of the agents' information is endogenous in the experimentation model of Bergemann and Välimäki (2010), in the procurement model of Krähmer and Strausz (2011), in the sponsored-search model of Kakade et al. (2013), in the bandit-auction model of Pavan, Segal, and Toikka (2014), in the matching model of Fershtman and Pavan (2016), and in the taxation model of Makris and Pavan (2016). This last paper is reviewed in Section 6, below; it considers a dynamic taxation problem in which the agents' productivity evolves endogenously as the result of learning-by-doing.

Related is also the literature on dynamic managerial compensation. Most of this literature studies optimal compensation schemes in a pure moral hazard setting (see, for example, Prendergast, 2002 for an earlier overview; Sannikov, 2013 for a more recent overview of the continuous-time contracting literature; and the references in Board, 2011 for the subset of this literature focusing on relational contracting). The part of this literature that is most related to the dynamic mechanism design literature is the one that assumes that the manager observes shocks to the cash flows prior to committing his

dynamic environments, see Aghion et al. (2012), Mueller (2015), and Penta (2015). Another strand of the literature studies screening and moral hazard problems in settings in which the principal lacks information about the type distribution, the set of available effort choices, or the technology used by nature to perturb the agent's action. See, for example, Segal (2003), Frankel (2012), Chassang (2013), Garrett (2014), Carroll (2015), and the references therein.

6 See also Krähmer and Strausz (2016) for a discussion of how the analysis of sequential screening in Courty and Li (2000) can be reconducted to a static screening problem with stochastic allocations.

effort (as in the taxation and in the regulation literature); see, for example, Edmans and Gabaix (2011), Edmans et al. (2012), Garrett and Pavan (2012), and Carroll and Meng (2016). This timing is also the one considered in the variational-approach paper by Garrett and Pavan (2015) reviewed in Section 5, below.

Most of the analysis in the dynamic mechanism design literature is in discrete time. One of the earlier papers in continuous time is Williams (2011). For a discussion of the developments of the continuous-time dynamic adverse selection literature and its connection to discrete time, see the recent paper by Bergemann and Strack (2015a) and the references therein.[7]

The dynamic mechanism design literature typically assumes that the designer can commit to her mechanism, with the dynamics of the allocations originating either in evolving private information or in the stochastic arrival and departure of goods and agents over time. A related literature on dynamic contracting under limited commitment investigates the dynamics of allocations in models in which the agents' private information is static but where the principal is unable to commit to future decisions. For earlier contributions to this literature, see, for example, Laffont and Tirole (1988), and Hart and Tirole (1988). For more recent contributions, see Skreta (2006, 2015), Battaglini (2007), Galperti (2015), Maestri (2016), Gerardi and Maestri (2016), Liu et al. (2015), Strulovici (2016), and the references therein. A particular form of limited commitment is considered in Deb and Said (2015). In that paper, the seller can commit to the dynamic contract she offers to each agent, but cannot commit to the contracts she offers to agents arriving in future periods. Partial commitment is also the focus of a recent paper in continuous time by Miao and Zhang (2015), in which both the principal and the agent can walk away from the relationship at any point in time after observing the evolution of the agent's income process.

Another assumption typically maintained in the dynamic mechanism design literature is that transfers can be used to incentivize the agents to report their private information (and/or to exert effort). A few papers investigate dynamic incentives in settings with or without evolving private information, in which transfers are not feasible. An early contribution to this literature is Hylland and Zeckhauser (1979). More recent contributions include Abdulkadiroğlu and Loertscher (2007), Miralles (2012), Kováč et al. (2014), Johnson (2015), Li et al. (2015), Frankel (2016), Johnson (2015), and Guo and Hörner (2016).

Related is also the literature on information design. For a survey of earlier contributions see Bergemann and Välimäki (2006). For more recent developments, including dynamic extensions, see Gershkov and Szentes (2009), Rayo and Segal (2010), Kamenica and Gentzkow (2011), Gentzkow and Kamenica (2015), Bergemann and Morris (2016), Ely et al. (2016), Doval and Ely (2016), Ely, Garrett, and Hinnosaar (2016), and the references

[7] See also Prat and Jovanovic (2014), Strulovici and Szydlowski (2015), and Williams (2015) for recent contributions.

therein. Hörner and Skrzypacz (2016) offer a useful survey of these recent developments. The canonical persuasion model assumes that the designer (the sender) can choose the information structure for the receiver at no cost. In contrast, Calzolari and Pavan (2006a, 2006b), consider models in which a principal first screens the private information of one, or multiple agents, and then passes a garbled version of this information to other agents, or other principals. The design of optimal disclosure rules in screening environments is also the focus of Bergemann and Pesendorfer (2007), Eső and Szentes (2007), Bergemann and Wambach (2015), and Nikandrova and Pancs (2015); all these papers study the design of optimal information structures in auctions.

Finally, dynamic mechanism design is related to the literature on information acquisition in mechanism design (see Bergemann and Välimäki (2002, 2006) and the references therein for earlier contributions, and Gershkov and Szentes (2009), and Krähmer and Strausz (2011) for some recent developments).

## 2 SIMPLE DYNAMIC SCREENING MODEL

In this section, I introduce a simple dynamic screening model that I use in the next four sections to illustrate some of the key ideas in the dynamic mechanism design literature.

The principal is a seller, the agent is a buyer. Their relationship lasts for $T \in \mathbb{N} \cup \{\infty\}$ periods, where $T$ can be either finite or infinite. Time is discrete and indexed by $t = 1, 2, \ldots, T$. 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 $q_t$.[8] The function $C(\cdot)$ 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, with $F_1$ denoting the initial distribution over $\Theta_1$ and, for all $t \geq 2$, $F_t(\cdot \mid \theta_{t-1})$ denoting the cdf of $\theta_t$ given $\theta_{t-1}$. Note that the above specification assumes the process is Markov and exogenous.

The sequence of events is the following.

- At $t = 0$, i.e., prior to entering any negotiations with the principal, the buyer privately learns $\theta_1$.

[8] The results for a static relationship can be read from the formulas below for the dynamic environment by setting $\delta = 0$.

- At $t = 1$, the seller offers a mechanism $\varphi = (\mathcal{M}, \phi)$. The latter consists of a collection of mappings

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

specifying a quantity–price pair for each possible history of messages $m^t \equiv (m_1, \ldots, m_t) \in \mathcal{M}_1 \times \cdots \times \mathcal{M}_t$, with $\mathcal{M} \equiv (\mathcal{M})_{t=1}^T$ and $\phi \equiv (\phi_t)_{t=1}^T$. A mechanism is thus equivalent to a menu of long-term contracts. If the buyer refuses to participate in $\varphi$, the game ends and both players obtain a payoff equal to zero. If the buyer chooses to participate in $\varphi$, he sends a message $m_1 \in \mathcal{M}_1$, receives quantity $q_1(m_1)$, pays a transfer $p_1(m_1)$, and the game moves to period 2.
- At the beginning of each period $t \geq 2$, the buyer privately learns $\theta_t$. 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 $t + 1$.
- $\ldots$
- At $t = T + 1$ the game is over (in case $T$ is finite).

**Remark** The game described above assumes that the principal (here the seller) perfectly commits to the mechanism $\varphi$. It also assumes that at any period $t \geq 2$ the buyer is constrained to stay in the relationship if he signed on in period 1. When the agent has "deep pockets," there are, however, simple ways to distribute the payments over time so that it is in the interest of the buyer to remain in the relationship at all periods, irrespective of what he did in the past.[9] ||

The principal's problem consists in designing a mechanism that disciplines the provision of quantity and the payments over time. Because the principal can commit, the Revelation Principle[10] applies and one can without loss of optimality restrict attention to direct mechanisms in which $\mathcal{M}_t = \Theta_t$ all $t$ and such that the agent finds it optimal to report truthfully at all periods. 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 it induces, where, for any $t \geq 1$, $q_t : \Theta^t \to \mathcal{Q}$ is the period-$t$ output policy and $p_t : \Theta^t \to \mathbb{R}$ the payment policy, with $\Theta^t = \Theta_1 \times \cdots \times \Theta_t$.[11] The principal designs $\chi$ so as to maximize

$$\mathbb{E}\left[\sum_t \delta^{t-1}(p_t(\theta^t) - C(q_t(\theta^t)))\right]$$

[9] See also Krähmer and Strausz (2015a) for a discussion of interim vs ex-post participation constraints in sequential secreening models.

[10] See, among others, Gibbard (1977), Green and Laffont (1979), Myerson (1979, 1986).

[11] A similar notation will be used hereafter to denote sequences of sets. For example, $\mathcal{A}^t = \mathcal{A}_1 \times \cdots \times \mathcal{A}_t$ with generic element $a^t = (a_1, \ldots, a_t)$.

subject to

$$\mathbb{E}\left[\sum_t \delta^{t-1}(\theta_t q_t(\theta^t) - p_t(\theta^t)) \mid \theta_1\right] \geq 0 \text{ for all } \theta_1 \in \Theta \tag{IR-1}$$

$$\mathbb{E}\left[\sum_{s\geq t} \delta^{s-t}(\theta_s q_s(\theta^s) - p_s(\theta^s)) \mid \theta^t\right]$$
$$\geq \mathbb{E}\left[\sum_{s\geq t} \delta^{s-t}(\theta_s q_s^\sigma(\theta^s) - p_s^\sigma(\theta^s)) \mid \theta^t\right] \text{ for all } t,\ \theta^t \in \Theta^t,\ \text{and } \sigma, \tag{IC-t}$$

where $\boldsymbol{\sigma}$ denotes an arbitrary continuation strategy for the game that starts in period $t$ after the agent has reported truthfully at all previous periods. For any $s \geq t$, any $\theta^s \in \Theta^s$, $q_s^\sigma : \Theta^s \to \mathcal{Q}$ and $p_s^\sigma : \Theta^s \to \mathbb{R}$ are the state-contingent policies induced by the continuation strategy $\sigma$ under the mechanism $\chi = \langle \mathbf{q}, \mathbf{p} \rangle$. Hereafter, unless otherwise specified, the expectation operator $\mathbb{E}[\cdot]$ is with respect to the entire type sequence $(\theta_s)_{s=1}^T$ under the kernels $F$.

Note that the above constraints require that the buyer finds it optimal to participate in period 1 and report truthfully "on path," i.e., conditional on having reported truthfully in previous periods. Because the environment is Markov, holding fixed the agent's reports at each period $s < t$, the agent's incentives in period $t \geq 2$ are invariant in the agent's true type $\theta_s$, $s < t$. Hence, the above (IC-t) constraints guarantee that the agent finds it optimal to report truthfully at all histories, not just "on-path."

## 3 TWO-PERIOD DISCRETE EXAMPLE

To illustrate the trade-offs that determine the dynamics of allocations under optimal mechanisms in the simplest possible way, consider the following environment in which $T = 2$, $\Theta_1 \equiv \{\bar{\theta}, \underline{\theta}\}$, $\underline{\theta} > 0$, $\Delta\theta \equiv \bar{\theta} - \underline{\theta} > 0$, and $\Theta_2 \equiv \{\underline{\theta} - \Delta\theta, \underline{\theta}, \bar{\theta}, \bar{\theta} + \Delta\theta\}$. That is, the buyer has either a high or a low valuation in period 1. In period 2, he then experiences a shock that either raises his valuation by $\Delta\theta$, leaves his valuation unchanged, or reduces his valuation by $\Delta\theta$. To simplify, also assume that the principal's production cost is quadratic, with $C(q) = q^2/2$, all $q$.

The probability that the buyer is a high type in period 1 (equivalently, the proportion of high types in the cross-section of the population) is $\Pr(\theta_1 = \bar{\theta}) = \nu$. Conditional on $\theta_1$, the transition probabilities are as follows: $\Pr(\bar{\theta} + \Delta\theta | \bar{\theta}) = \bar{x}$, $\Pr(\bar{\theta}|\bar{\theta}) = \bar{\alpha}$, $\Pr(\underline{\theta}|\bar{\theta}) = 1 - \bar{x} - \bar{\alpha}$, $\Pr(\bar{\theta}|\underline{\theta}) = \underline{x}$, $\Pr(\underline{\theta}|\underline{\theta}) = \underline{\alpha}$, and $\Pr(\underline{\theta} - \Delta\theta|\underline{\theta}) = 1 - \underline{x} - \underline{\alpha}$. Figure 1 illustrates the situation under consideration.

As usual, the game is solved backwards. Let $U_2^A(\theta_2; \hat{\theta}_1, \hat{\theta}_2) \equiv \theta_2 q_2(\hat{\theta}_1, \hat{\theta}_2) - p_2(\hat{\theta}_1, \hat{\theta}_2)$ denote the agent's period-2 flow payoff when the true period-2 type is $\theta_2$ and the agents reported $(\hat{\theta}_1, \hat{\theta}_2)$ in the two periods. Note

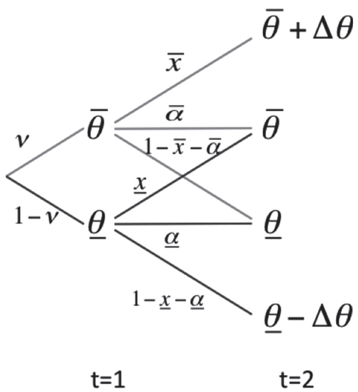


Figure 1 Evolution of Agent's Type.

that the flow period-2 payoff $U_2^A(\theta_2; \hat{\theta}_1, \hat{\theta}_2)$ does not depend on the agent's true period-1 type, $\theta_1$.

Next, let $V_2^A(\theta_1, \theta_2) \equiv U_2^A(\theta_2; \theta_1, \theta_2)$ denote the agent's period-2 flow payoff under truthful reporting (because it is irrelevant whether or not the period-1 report coincides with the true period-1 type, I am replacing $\hat{\theta}_1$ with $\theta_1$ to facilitate the notation).

Incentive compatibility in period 2 then requires that, for all $\theta_1 \in \Theta_1$, all $\theta_2, \hat{\theta}_2 \in \Theta_2$,

$$V_2^A(\theta_1, \theta_2) \geq U_2^A(\theta_2; \theta_1, \hat{\theta}_2).$$

Because the flow payoffs $\theta_t q_t - p_t$ satisfy the increasing differences property, it is well known, from static mechanism design, that incentive compatibility in period 2 requires that, for all $\theta_1 \in \Theta_1$, the output schedules $q_2(\theta_1, \cdot)$ be nondecreasing in $\theta_2$ and the payments $p_2(\theta_1, \cdot)$ satisfy the following conditions:

$$\Delta\theta q_2\left(\theta_1, \bar{\theta}\right) \leq V_2^A(\theta_1, \bar{\theta}+\Delta\theta) - V_2^A(\theta_1, \bar{\theta}) \leq \Delta\theta q_2\left(\theta_1, \bar{\theta}+\Delta\theta\right) \tag{1}$$

$$\Delta\theta q_2\left(\theta_1, \underline{\theta}\right) \leq V_2^A(\theta_1, \bar{\theta}) - V_2^A(\theta_1, \underline{\theta}) \leq \Delta\theta q_2\left(\theta_1, \bar{\theta}\right) \tag{2}$$

$$\Delta\theta q_2\left(\theta_1, \underline{\theta}-\Delta\theta\right) \leq V_2^A(\theta_1, \underline{\theta}) - V_2^A(\theta_1, \underline{\theta}-\Delta\theta) \leq \Delta\theta q_2\left(\theta_1, \underline{\theta}\right), \tag{3}$$

both for $\theta_1 = \bar{\theta}$ and for $\theta_1 = \underline{\theta}$. Along with the monotonicity of the output schedules $q_2(\theta_1, \cdot)$, the above constraints are not only necessary but also sufficient for period-2 incentive compatibility.

Next, let

$$U_1^A(\theta_1; \hat{\theta}_1) \equiv \theta_1 q_1(\hat{\theta}_1) - p_1(\hat{\theta}_1) + \delta\mathbb{E}\left[V_2^A(\hat{\theta}_1, \theta_2) \mid \theta_1\right]$$

denote the payoff that a buyer with initial type $\theta_1$ expects from reporting $\hat{\theta}_1$ in period 1 and, then, reporting truthfully at $t = 2$. Observe that the same period-2 incentive-compatibility constraints (1)–(3) that guarantee that it is optimal for the buyer to report truthfully in period 2 after having reported truthfully in period 1 also guarantee that it is optimal for him to report truthfully after having lied in period 1. Also, note that the expectation in the above definition of $U_1^A(\theta_1; \hat{\theta}_1)$ is over $\theta_2$ given the true period-1 type $\theta_1$. Finally, let $V_1^A(\theta_1) = U_1^A(\theta_1; \theta_1)$ denote the payoff that the buyer expects over the entire relationship under a truthful strategy in both periods, given his period-1 type $\theta_1$. The following is then a necessary condition for incentive-compatibility at $t = 1$:

$$\begin{aligned} V_1^A(\bar{\theta}) \geq V_1^A(\underline{\theta}) + \Delta\theta q_1(\underline{\theta}) + \delta\{\bar{x} V_2^A(\underline{\theta}, \bar{\theta} + \Delta\theta) \\ + \left(\bar{\alpha} - \underline{x}\right) V_2^A(\underline{\theta}, \bar{\theta}) + \left(1 - \bar{\alpha} - \bar{x} - \underline{\alpha}\right) V_2^A(\underline{\theta}, \underline{\theta}) \\ - \left(1 - \underline{\alpha} - \underline{x}\right) V_2^A(\underline{\theta}, \underline{\theta} - \Delta\theta)\}. \end{aligned} \tag{4}$$

Importantly, note that the right-hand side of (4) is written so as to highlight the extra payoff (or, equivalently, the informational rent) that a period-1 high type enjoys relative to a period-1 low type. The term $\Delta\theta q_1(\underline{\theta})$ is the familiar term from static mechanism design capturing the extra utility that the high type derives from the period-1 quantity $q_1(\underline{\theta})$ supplied to the period-1 low type. The term in curly brackets is the novel effect specific to the dynamic setting. It captures the discounted expectation of the second-period continuation payoff $V_2^A(\underline{\theta}_1, \theta_2)$ that follows the reporting of a low type in period 1, but where the discounting combines the time factor, $\delta$, with terms that capture the differences in the transition probabilities between the period-1 high and low types.

Now observe that the period-1 IC constraint (4) can be conveniently rewritten as

$$\begin{aligned} V_1^A(\bar{\theta}) \geq V_1^A(\underline{\theta}) + \Delta\theta q_1(\underline{\theta}) + \delta \Big\{\bar{x} \left[V_2^A(\underline{\theta}, \bar{\theta} + \Delta\theta) - V_2^A(\underline{\theta}, \bar{\theta})\right] \\ + \left(\bar{x} + \bar{\alpha} - \underline{x}\right) \left[V_2^A(\underline{\theta}, \bar{\theta}) - V_2^A(\underline{\theta}, \underline{\theta})\right] \\ + \left(1 - \underline{\alpha} - \underline{x}\right) \left[V_2^A(\underline{\theta}, \underline{\theta}) - V_2^A(\underline{\theta}, \underline{\theta} - \Delta\theta)\right]\Big\}. \end{aligned} \tag{5}$$

It is then evident from (5) that it is not the level of the continuation payoffs $V_2^A(\underline{\theta}, \theta_2)$ that determines the expected surplus that the high type can guarantee for himself over and above the one of the low type, but the “rate” by which such continuation payoffs change with the period-2 type $\theta_2$. Also, note that the terms in the round brackets in (5) are the differences in the *survival rates* across the two period-1 types. That is,

$$\begin{aligned} \bar{x} &= \Pr(\theta_2 \geq \bar{\theta} + \Delta\theta|\bar{\theta}) - \Pr(\theta_2 \geq \bar{\theta} + \Delta\theta|\underline{\theta}) \\ \bar{x} + \bar{\alpha} - \underline{x} &= \Pr(\theta_2 \geq \bar{\theta}|\bar{\theta}) - \Pr(\theta_2 \geq \bar{\theta}|\underline{\theta}) \\ 1 - \underline{\alpha} - \underline{x} &= \Pr(\theta_2 \geq \underline{\theta}|\bar{\theta}) - \Pr(\theta_2 \geq \underline{\theta}|\underline{\theta}). \end{aligned}$$

To reduce the informational rent of the period-1 high type, the principal can then either distort downwards the period-1 output provided to the period-1 low type, $q_1(\underline{\theta})$, or the various differences in the period-2 continuation payoffs across adjacent period-2 types, following a report $\underline{\theta}$ of a low type in period 1. Incentive compatibility in period 2, however, imposes lower bounds on such differentials, as one can see from (1)–(3). Consider then a *relaxed program* in which the following constraints are neglected: (a) the period-1 participation constraint of the high type; (b) the period-1 incentive-compatibility constraint of the low type; (c) all period-2 incentive compatibility constraints in the contract for the period-1 high type (that is, (1)–(3) for $\theta_1 = \bar{\theta}$); (d) all period-2 upward adjacent incentive-compatibility constraints in the contract for the period-1 low type (i.e., the right-hand inequalities in the constraints (1)–(3) for $\theta_1 = \underline{\theta}$). Formally, the relaxed program can be stated as follows:

$$\mathcal{P}_r : \begin{cases} \max_{\chi=\langle \mathbf{q},\mathbf{p}\rangle} \mathbb{E}\left\{\theta_1 q_1(\theta_1) - C(q_1(\theta_1)) + \delta\left[\theta_2 q_2(\theta_1,\theta_2) - C(q_2(\theta_1,\theta_2))\right] \right. \\ \qquad\qquad \left. - V_1^A(\theta_1)\right\} \\ \text{subject to } V_1^A(\underline{\theta}) \geq 0, \ \text{(5) and} \\ V_2^A(\underline{\theta},\bar{\theta}+\Delta\theta) - V_2^A(\underline{\theta},\bar{\theta}) \geq \Delta\theta q_2\left(\underline{\theta},\bar{\theta}\right) \\ V_2^A(\underline{\theta},\bar{\theta}) - V_2^A(\underline{\theta},\underline{\theta}) \geq \Delta\theta q_2\left(\underline{\theta},\underline{\theta}\right) \\ V_2^A(\underline{\theta},\underline{\theta}) - V_2^A(\underline{\theta},\underline{\theta}-\Delta\theta) \geq \Delta\theta q_2\left(\underline{\theta},\underline{\theta}-\Delta\theta\right). \end{cases}$$

Now suppose the process is *stochastically monotone* (in the sense that the distribution of $\theta_2$ given $\theta_1 = \bar{\theta}$ first-order-stochastic dominates the distribution of $\theta_2$ given $\theta_1 = \underline{\theta}$), which is the case if and only if $\bar{x} + \bar{\alpha} - \underline{x} \geq 0$. The following result is then true:

**Proposition 1** *Suppose the process is* stochastically monotone. *Then any solution to the relaxed program is such that all constraints in $\mathcal{P}_r$ bind, and is characterized by the following output schedules:*[12]

$$\begin{aligned} q_1(\overline{\theta}) &= q_1^{FB}(\overline{\theta}); \\ q_2(\overline{\theta},\theta_2) &= q_2^{FB}(\overline{\theta},\theta_2), \ \forall \theta_2 \in \Theta_2; \\ q_1(\underline{\theta}) &= \max\{q_1^{FB}(\underline{\theta}) - \tfrac{\nu}{1-\nu}\Delta\theta; 0\}; \\ q_2(\underline{\theta},\bar{\theta}) &= \max\{q_2^{FB}(\underline{\theta},\bar{\theta}) - \left(\tfrac{\nu}{1-\nu}\right)\left(\tfrac{\bar{x}}{\underline{x}}\right)\Delta\theta; 0\}; \\ q_2(\underline{\theta},\underline{\theta}) &= \max\{q_2^{FB}(\underline{\theta},\underline{\theta}) - \left(\tfrac{\nu}{1-\nu}\right)\left(\tfrac{\bar{x}+\bar{\alpha}-\underline{x}}{\underline{\alpha}}\right)\Delta\theta; 0\}; \\ q_2(\underline{\theta},\underline{\theta}-\Delta\theta) &= \max\{q_2^{FB}(\underline{\theta},\underline{\theta}-\Delta\theta) - \left(\tfrac{\nu}{1-\nu}\right)\Delta\theta; 0\}, \end{aligned}$$

*where $q_1^{FB}(\theta_1) = \theta_1$, and $q_2^{FB}(\theta_1,\theta_2) = \theta_2$ are the first-best output schedules (i.e., the schedules that would be implemented under complete information). The solution to the relaxed program $\mathcal{P}_r$ coincides with the solution to the full program if and only if the schedule $q_2(\underline{\theta},\cdot)$ is nondecreasing in $\theta_2$.*

[12] Not surprisingly, the output schedules that solve the relaxed and the full programs are not uniquely defined at $(\overline{\theta},\underline{\theta}-\Delta\theta)$ and $(\underline{\theta},\bar{\theta}+\Delta\theta)$, for these histories have zero measure.

Hence, the solution to the relaxed program entails no distortions in either period in the output provided to the period-1 high type and downward distortions in either period in the output provided to the period-1 low type. Importantly, the distortions in the output provided to the period-1 low type need not decrease over time or be monotone in the period-2 type.

**Corollary 1** *The output schedule $q_2(\underline{\theta}, \cdot)$ that solves the relaxed program is nondecreasing in $\theta_2$ (and hence part of an optimal contract) if and only if*

$$-\frac{\nu}{1-\nu} \leq 1 - \left(\frac{\nu}{1-\nu}\right)\left(\frac{\bar{x} + \bar{\alpha} - \underline{x}}{\underline{\alpha}}\right) \leq 2 - \left(\frac{\nu}{1-\nu}\right)\left(\frac{\bar{x}}{\underline{x}}\right).$$

*Distortions are weakly higher in period 2 than in period 1 if and only if*

$$1 \leq \frac{\bar{x} + \bar{\alpha} - \underline{x}}{\underline{\alpha}} \leq \frac{\bar{x}}{\underline{x}}.$$

*Because the above two sets of conditions are not mutually exclusive, distortions under optimal contracts need not be monotone in type or in time.*

The example also illustrates the importance of not restricting attention to oversimplified processes. Suppose, for example, that one were to assume that the agent's marginal value for the principal's product takes only two values, $\underline{\theta}$ or $\bar{\theta}$, in each period. In the context of our example, this amounts to imposing that $\bar{x} = 0 = 1 - \underline{\alpha} - \underline{x}$. In this case, $q_2(\underline{\theta}, \bar{\theta}) = q_2^{FB}(\underline{\theta}, \bar{\theta})$ whereas

$$\begin{aligned} q_1(\underline{\theta}) &= \max\{q_1^{FB}(\underline{\theta}) - \tfrac{\nu}{1-\nu}\Delta\theta; 0\}; \\ q_2(\underline{\theta}, \underline{\theta}) &= \max\{q_2^{FB}(\underline{\theta}, \underline{\theta}) - \left(\tfrac{\nu}{1-\nu}\right)\left(\tfrac{\bar{\alpha}+\underline{\alpha}-1}{\underline{\alpha}}\right)\Delta\theta; 0\}; \\ q_2(\underline{\theta}, \bar{\theta}) &= q_2^{FB}(\underline{\theta}, \bar{\theta}). \end{aligned} \tag{6}$$

In this case, as soon as the agent's type turns high, he receives first-best output. Furthermore, distortions in the output provided to the the period-1 low type decrease monotonically over time. As shown in Battaglini (2005), these properties extend to arbitrary horizons, as long as the process takes only two values, and remains stochastically monotone.

A central question for the dynamic mechanism design literature is then what properties of the process governing the evolution of the agents' private information are responsible for the dynamics of the distortions under optimal contracts. A conjecture in the earlier literature is that distortions should decline over time when the correlation between the initial type, $\theta_1$, and the subsequent ones, $\theta_t$, declines over time. To see the problem with this conjecture, consider the same example above in which the type process takes only two values in each period. However, assume now that $\bar{\theta}$ is an absorbing state so that $\bar{\alpha} = 1$. It is easy to see from the formulas in (6) that, while in this case distortions decrease in expectation over time, they remain constant as long as the agent's type remains low. Furthermore, even the property that distortions decline in expectation should not be taken for granted. Suppose, for example that the agent's type follows a random walk. In the example above, this amounts to assuming that $\theta_2 = \theta_1 + \varepsilon$ with $\varepsilon \in \{-\Delta\theta, 0, +\Delta\theta\}$ and with the distribution over $\varepsilon$ independent of $\theta_1$ so that $\bar{x} = \underline{x}$ and $\bar{\alpha} = \underline{\alpha}$. Then use the formulas in Proposition 1 to verify that the distortions in the contract of the period-1 low type remain constant over time. As I discuss further in the next section, this property extends to arbitrary horizons and to general random walk processes. Hence, despite the fact that, when the process is a random walk, the correlation between the agent's period-1 type and his period-$t$ type declines with $t$, distortions remain constant over time.

Another observation that the example offers is that the volatility of the shocks, or their persistence, need not matter for the dynamics of distortions. To see this, suppose that $\bar{\alpha} = \underline{\alpha} = 1$, in which case the agent's type is constant over time. Clearly, this is a special case of the random walk case discussed above, where the volatility of the period-2 shocks is zero. Then, again, distortions are constant over time. In fact, as shown first in Baron and Besanko (1984), when types are constant over time and payoffs are time-additively-separable, as assumed here, the optimal allocations can be implemented by offering the same static contract in each period.

Another polar case of interest is when types are independent over time. In the context of this simple example, this amounts to assuming that $\bar{x} = 1 - \underline{\alpha} - \underline{x} = 0$ and that $\bar{\alpha} = \underline{x}$. In this case there are no distortions in the second period allocation, irrespective of the first-period type.

It should be clear from the discussion above that, as in static problems, distortions in optimal contracts are driven by the familiar trade-off between efficiency and rent-extraction, *evaluated from period 1's perspective*. That is, from the perspective of the period at which the participation constraints bind. It should also be clear that neither the correlation between the initial type and his future types, nor the agent's ability to forecast his future types (as captured, for example, by the inverse of the volatility of the forecast error) is what determines the way the principal solves the above trade-off intertemporally. The question is then what specific properties of the process governing the evolution of the agents' private information are responsible for the dynamics of distortions under optimal contracts? I will provide a formal answer to this question in the next section, after developing an approach that permits one to extend the analysis to a broader class of processes (and of contracting problems). Before doing so, it is, however, instructive to rewrite (5) as follows:

$$\frac{V_1^A(\bar{\theta}) - V_1^A(\underline{\theta})}{\Delta\theta} = q_1(\underline{\theta}) + \delta\mathbb{E}\left[I_2(\underline{\theta}, \theta_2) \cdot q_2\left(\underline{\theta}, \theta_2\right) \mid \theta_1 = \underline{\theta}\right], \quad (7)$$

where

$$I_2(\underline{\theta}, \theta_2) = \frac{\Pr(\tilde{\theta}_2 > \theta_2|\bar{\theta}) - \Pr(\tilde{\theta}_2 > \theta_2|\underline{\theta})}{\Pr(\tilde{\theta}_2 = \theta_2|\underline{\theta})}. \quad (8)$$

The function $I_2(\underline{\theta}, \theta_2)$ in (8) captures the different probability the two period-1 types assign to having a period-2 type above $\theta_2$, normalized by the probability that the period-1 low type assigns to having a period-2 type equal to $\theta_2$. As it will become clear below, these functions are the discrete-type analogs of what Pavan, Segal, and Toikka (2014) refer to as "*impulse response functions*."

As in static mechanism design, the surplus the principal must grant to the period-1 high type to induce him to report truthfully originates from the surplus that this type can guarantee himself by mimicking the period-1 low type. Contrary to static mechanism design, though, such surplus combines the difference in the period-1 utility of consuming the quantity $q_1(\underline{\theta})$ sold to the low type with the difference in the continuation payoffs. The latter are in turn determined by the quantities supplied in the second period to the period-1 low type (which are responsible for the period-2 informational rents) scaled by the different probabilities the two period-1 types assign to being able to enjoy such rents, as captured by the impulse response functions.

## 4 MYERSONIAN APPROACH

The example in the previous section illustrates some of the key ideas in the design of optimal dynamic mechanisms. It also shows the importance of not over-simplifying the model when it comes to the predictions the theory delivers for the dynamics of the distortions under optimal contracts. In particular, the example shows the importance of accommodating for more than two types. As in other contexts, working with a large but finite number of types is more tedious than working with a continuum of types. The literature has thus investigated ways of extending the analysis of mechanism design with a continuum of types to dynamic settings. The approach is similar to the one pioneered by Myerson (1981), but adapted to account for the gradual resolution of uncertainty and for the possibility of a stream of decisions. I revisit some of the key results of this approach in this section.

There are a few difficulties in extending the Myersonian approach to dynamic settings. The first is in identifying primitive properties (on payoffs and type processes) under which it is a necessary condition for incentive compatibility that each agent's equilibrium payoff satisfies, at each period, an *envelope formula* analogous to the familiar one

$$V^A(\theta) = V^A(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} q(s)ds \tag{9}$$

for static environments.

A second difficulty is in identifying appropriate *dynamic monotonicity conditions* on the allocation rule that, when paired with the envelope formulas for the equilibrium payoffs, guarantee that the agents have the incentives to report truthfully at all histories.

The difficulty in identifying primitive conditions implying that the envelope representation of the equilibrium payoffs is a necessary condition for incentive compatibility comes from the fact that (a) the policies in the mechanism cannot be assumed to be "smooth," for this may affect the characterization of the optimal mechanisms, and (b) one cannot use directly the conditions justifying the envelope theorem in static models (e.g., Milgrom and Segal, 2002), for the latter require each agent's payoff to be smooth in the agent's type *across all possible strategies* and, in dynamic settings, such a property cannot be guaranteed, at least not for arbitrary mechanisms.

Importantly, these difficulties are conceptual, not technical. In particular, the difficulty in guaranteeing that each agent's payoff is smooth in the initial type, across all possible continuation strategies, is akin to the difficulty in establishing revenue equivalence in multiagent static settings with correlated information. Similarly, the difficulty in identifying appropriate dynamic monotonicity conditions is akin to the difficulty in identifying minimally sufficient conditions in static screening problems with multidimensional types and allocations.

In this section, I describe the approach in Pavan, Segal, and Toikka (2014) to address such difficulties. The exposition here is informal. I refer the reader to the original paper for details.

## 4.1 Incentive Compatibility, Envelopes, and Payoff Equivalence

In order to identify primitive conditions (on payoffs and type processes) implying that the equilibrium payoffs must satisfy an envelope representation akin to the one in static models, across all possible mechanisms, Pavan, Segal, and Toikka (2014) use a state representation of the evolution of the agents' private information similar to the one in Eső and Szentes (2007) – see also Eső and Szentes (2015).[13] That is, let $\theta_t = Z_t(\theta_1, \varepsilon)$ where $\varepsilon$ is a vector of random variables, independent of $\theta_1$, drawn from some distribution $G$ over $\mathcal{E} \equiv \prod_{t=1}^{T} \mathcal{E}_t$, with each $\mathcal{E}_t \subset \mathbb{R}$. Any stochastic process can be described this way. For example, given the kernels $F \equiv (F_t)$, for any $t \geq 2$, one can let $F_t^{-1}(\varepsilon_t|\theta_{t-1}) \equiv \inf\{\theta_t : F_t(\theta_t|\theta_{t-1}) \geq \varepsilon_t\}$ with each $\varepsilon_t$ drawn from a Uniform distribution over $(0, 1)$, independently of $\theta_1$ and of any $\varepsilon_s$, $s \neq t$. The random variable $F_t^{-1}(\varepsilon_t|\theta_{t-1})$ is then distributed according to the cdf $F_t(\cdot|\theta_{t-1})$. One can then let $Z_t(\theta_1, \varepsilon) = F_t^{-1}(\varepsilon_t|Z_{t-1}(\theta_1, \varepsilon))$, with $Z^{t-1}(\theta_1, \varepsilon) \equiv (Z_\tau(\theta_1, \varepsilon))_{\tau=1}^{t-1}$ constructed inductively starting from $Z_1(\theta_1) \equiv \theta_1$.[14] The above representation is referred to as the *canonical representation* of $F$ in Pavan, Segal, and Toikka (2014). While the canonical representation is a valid representation for all processes, it is sometimes convenient to work with alternative state representations. For example, suppose that $\theta_t$ evolves according to an AR(1) process, so that $\theta_t = \gamma\theta_{t-1} + \varepsilon_t$. Then a natural state representation is

$$\theta_t = Z_t(\theta_1, \varepsilon) = \gamma^{t-1}\theta_1 + \gamma^{t-2}\varepsilon_2 + \cdots + \gamma\varepsilon_{t-1} + \varepsilon_t.$$

[13] These papers use the state representation as a tool to study optimal disclosure policies as opposed to a tool to identify primitive conditions validating the envelope theorem.

[14] This construction is standard; see the second proof of Kolmogorov's extension theorem in Billingsley (1995, p.490).

A similar state representation can be used to describe the agent's private information at any period $t$ as a function of his period-$s$ type history $\theta^s \equiv (\theta_1, \ldots, \theta_s)$, $s < t$, and subsequent shocks. Thus, for any $s \geq 1$, any $t > s$, let $Z_{(s),t}(\theta^s, \varepsilon)$ denote the representation of $\theta_t$ as a function of $\theta^s$ and independent shocks $\varepsilon$ that, along with $\theta^s$, are responsible for $\theta_t$. Hereafter, the index $s$ is dropped from the subscripts when $s = 1$ to ease the exposition (that is, for any $t$, any $(\theta_1, \varepsilon)$, $Z_t(\theta_1, \varepsilon) \equiv Z_{(1),t}(\theta_1, \varepsilon)$).

Next, let $\|\cdot\|$ denote the discounted L-1 norm on $\mathbb{R}^\infty$ defined by $\|y\| \equiv \sum_{t=0}^{\infty} \delta^t |y_t|$.

**Definition 1** *The process is* regular *if, for any $s \geq 1$, there exist a function $C_{(s)} : \mathcal{E} \to \mathbb{R}^\infty$ satisfying $\mathbb{E}[||C_{(s)}(\varepsilon)||] \leq B$ for some constant $B$ independent of $s$, such that for all $t \geq s$, $\theta^s \in \Theta^s$, and $\varepsilon \in \mathcal{E}$, $Z_{(s),t}(\theta^s, \varepsilon)$ is a differentiable function of $\theta_s$ with $\left|\partial Z_{(s),t}(\theta^s, \varepsilon)/\partial \theta_s\right| \leq C_{(s),t-s}(\varepsilon)$.*[15]

In words, a process is regular if it admits a state representation where the $Z_{(s),t}$ functions are differentiable with bounded derivatives (but where the bounds are allowed to vary with the state).

Given the above notation, the *impulse response* of $\theta_t$ to $\theta_s$ is then defined as follows:[16]

$$I_{(s),t}(\theta^t) \equiv \mathbb{E}\left[\frac{\partial Z_{(s),t}(\theta^s, \varepsilon)}{\partial \theta_s} \mid Z_{(s),t}(\theta^s, \varepsilon) = \theta^t\right],$$

where the expectation is over the shocks $\varepsilon$ given $\theta^s$. Again, when $s = 1$, to simplify the notation, drop the index $(s)$ from the subscripts so that by $I_t(\theta^t) \equiv I_{(1),t}(\theta^t)$ all $t > 1$ all $\theta^t$. Finally, for all $s \geq 1$, all $\theta^s$, $I_{(s),s}(\theta^s) \equiv 1$. For example, when $\theta_t$ follows an AR(1) process, the impulse response of $\theta_t$ to $\theta_1$ is simply $I_t(\theta_1, \varepsilon) = \gamma^{t-1}$. More generally, the impulse response functions are themselves stochastic processes capturing how variations in types at any given period propagate through the entire process. For example, when

$$\theta_t = Z_t(\theta_1, \varepsilon) = \theta_1 \times \varepsilon_2 \times \cdots \times \varepsilon_t$$

the period-1 impulse response functions are given by $I_t(\theta_1, \varepsilon) = \varepsilon_2 \times \cdots \times \varepsilon_t$. When the process is Markov and the kernels $F_t(\theta_t|\theta_{t-1})$ are continuously differentiable in $(\theta_t, \theta_{t-1})$, the canonical representation introduced above yields

[15] For any $\varepsilon$, the term $C_{(s),t-s}(\varepsilon)$ is the $(t-s)$-compoment of the sequence $C_s(\varepsilon)$.

[16] Special cases of the impulse response functions (and of the ICFOC conditions below) have been identified by Baron and Besanko (1984), Besanko (1985), Courty and Li (2000), and Eső and Szentes (2007), among others. In particular, Baron and Besanko (1984) refer to these functions as to "the informativeness measure." For reasons discussed in detail in Pavan, Segal, and Toikka (2014), we favor the expression "impulse response functions" to highlight that the information that early types contain about future ones is not the precise property of the type process responsible for the dynamics of distortions under optimal contracts – recall the discussion in the previous section.

the following expression for the impulse responses:

$$I_{(s),t}(\theta^t) = \prod_{\tau=s+1}^{t} \left( -\frac{\partial F_\tau(\theta_\tau|\theta_{\tau-1})/\partial\theta_{\tau-1}}{f_\tau(\theta_\tau|\theta_{\tau-1})} \right).$$

Note that, when $s = 1$ and $t = 2$, the impulse response function of $\theta_2$ to $\theta_1$ is then given by

$$I_2(\theta_1, \theta_2) = -\frac{\partial F_2(\theta_2|\theta_1)/\partial\theta_1}{f_2(\theta_2|\theta_1)} = \frac{\partial[1 - F_2(\theta_2|\theta_1)]/\partial\theta_1}{f_2(\theta_2|\theta_1)},$$

which is the continuous-type analog of the discrete formula in (8). For the analogs of such formulas in continuous time, see Bergemann and Strack (2015a).

The advantage of the state representation when it comes to identifying primitive conditions under which the dynamic analog of the envelope representation of the equilibrium payoffs is a necessary condition for incentive compatibility is that it permits one to focus on a subset of possible continuation strategies for the agent that are indexed by state variables, the shocks, that are orthogonal to the initial private information. To see this, let $T = 2$ and assume the agent observes the shocks $\varepsilon$ in addition to his types (Notice that $(\theta_1, \theta_2)$ remains a sufficient statistic for $(\theta_1, \theta_2, \varepsilon)$ when it comes to the agent's payoff). Because there are only two periods, then, without loss of generality, $\varepsilon$ can be assumed to be unidimensional. Given any mechanism $\chi = \langle \mathbf{q}, \mathbf{p} \rangle$, then construct an auxiliary mechanism $\tilde{\chi} = \langle \tilde{\mathbf{q}}, \tilde{\mathbf{p}} \rangle$ in which the agent is asked to report $\theta_1$ in period 1 and the shock $\varepsilon$ in period 2. The mechanism $\tilde{\chi} = \langle \tilde{\mathbf{q}}, \tilde{\mathbf{p}} \rangle$ is obtained from $\chi = \langle \mathbf{q}, \mathbf{p} \rangle$ by letting $(\tilde{q}_1(\cdot), \tilde{p}_1(\cdot)) = (q_1(\cdot), p_1(\cdot))$ and then letting $\tilde{q}_2(\theta_1, \varepsilon) = q_2(\theta_1, Z_2(\theta_1, \varepsilon))$ and $\tilde{p}_2(\theta_1, \varepsilon) = p_2(\theta_1, Z_2(\theta_1, \varepsilon))$, where $Z_2(\theta_1, \varepsilon)$ is the function describing the state representation of the process introduced above. Then let

$$U^A(\theta_1; \hat{\theta}_1, \sigma) \equiv \theta_1 \tilde{q}_1(\hat{\theta}_1) - \tilde{p}_1(\hat{\theta}_1) + \delta \mathbb{E}\left[ Z_2(\theta_1, \varepsilon) \tilde{q}_2(\hat{\theta}_1, \sigma(\varepsilon)) - \tilde{p}_2(\hat{\theta}_1, \sigma(\varepsilon)) | \theta_1 \right]$$

denote the payoff that the initial type $\theta_1$ expects from reporting the message $\hat{\theta}_1$ in period 1 and then following the reporting strategy $\sigma(\varepsilon)$ in period 2, where $\sigma(\varepsilon)$ is the report the agent makes in period 2 when the period-2 shock is $\varepsilon$, and where the expectation is over the shocks $\varepsilon$, given $\theta_1$. Once the agent's behavior is indexed this way, fixing $(\hat{\theta}_1, \sigma)$, a variation in $\theta_1$ no longer triggers a variation in the agent's reports. This is convenient, for it bypasses the problem alluded to above of guaranteeing that the agent's payoff be differentiable in the true type $\theta_1$ for all possible reporting strategies.[17]

[17] Notice that if, instead, one were to index the agent's reporting plan by $(\hat{\theta}_1, \sigma)$ with the period-2 report $\sigma(\theta_2)$ now depending on $\theta_2$, as opposed to the orthogonalized innovation $\varepsilon$, then a variation in $\theta_1$, by triggering a variation in the period-2 report, would possibly lead to non-differentiability of $U^A(\theta_1; \hat{\theta}_1, \sigma)$ in $\theta_1$.

The construction above, in turn, can be used to validate the envelope theorem. Provided $U^A(\theta_1; \hat{\theta}_1, \sigma)$ is differentiable and equi-Lipschitz continuous in $\theta_1$, the value function $\sup_{(\hat{\theta}_1, \sigma)} U^A(\theta_1; \hat{\theta}_1, \sigma)$ is then Lipschitz continuous in $\theta_1$ (this follows from the same arguments as in Milgrom and Segal (2002)'s atemporal analysis).[18] Now, if the mechanism $\chi = \langle \mathbf{q}, \mathbf{p} \rangle$ is incentive compatible in the primitive environment, then reporting $\hat{\theta}_1 = \theta_1$ in period 1 and then reporting the shocks $\varepsilon$ truthfully in period 2 (that is, following the strategy $\sigma^{truth}(\varepsilon) = \varepsilon$) must be optimal for the agent in the auxiliary mechanism $\tilde{\chi} = \langle \tilde{\mathbf{q}}, \tilde{\mathbf{p}} \rangle$. The envelope theorem then implies that a necessary condition for incentive compatibility of $\chi = \langle \mathbf{q}, \mathbf{p} \rangle$ is that, under $\chi = \langle \mathbf{q}, \mathbf{p} \rangle$, $V_1^A(\theta_1)$ is Lipschitz continuous with derivative equal to

$$\frac{dV_1^A(\theta_1)}{d\theta_1} = q_1(\theta_1) + \delta \mathbb{E}\left[\frac{\partial Z_2(\theta_1, \varepsilon)}{\partial \theta_1} \tilde{q}_2(\theta_1, \varepsilon)\right]$$
$$= \mathbb{E}\left[\sum_{s=1}^{s=2} \delta^{s-1} I_s(\theta^s) q_s(\theta^s) \mid \theta_1\right], \tag{10}$$

where the last equality uses the definition of the impulse response functions along with the relation between the policies $\chi = \langle \mathbf{q}, \mathbf{p} \rangle$ and the policies $\tilde{\chi} = \langle \tilde{\mathbf{q}}, \tilde{\mathbf{p}} \rangle$. Note that the above envelope formula is the continuous-type analog of Condition (7) in the discrete case. As mentioned above, this formula is a necessary condition for incentive compatibility; it must hold in order for the agent to prefer reporting truthfully in both periods than lying in period 1 and then reporting the orthogonalized shocks truthfully. To obtain a complete characterization of incentive compatibility, one first extends the above envelope formula to arbitrary horizons and then combines it with appropriate monotonicity conditions on the allocation rule $\mathbf{q}$ guaranteeing that one-stage deviations from truthful reporting are suboptimal. When applied to a Markov environment (such as the one under consideration here), this approach yields a complete characterization of incentive compatibility, as shown in Theorem 1, below (for a general treatment in richer environments with an arbitrary number of agents and decision-controlled processes, see Pavan, Segal, and Toikka, 2014).[19]

[18] The differentiability and equi-Lipschitz continuity of $U^A$ in $\theta_1$ in turn can be guaranteed by assuming differentiability and equi-Lipschitz continuity of the $Z_2$ functions with appropriate bounds on the expected NPV of the quantity schedules. See Pavan et al. (2014) for details.

[19] Here, I assume that the principal is restricted to offering contracts such that, for any $t \geq 1$, any $\theta^t$, $\mathbb{E}\left[\sum_{s \geq t} \delta^{s-t} I_{(s),t}(\theta^s) q_s(\theta^s) \mid \theta^t\right] \leq K$ for some $K \in \mathbb{R}_{++}$ arbitrarily large. This restriction, which is stronger than needed, is used to validate the representation of the equilibrium payoffs in the theorem below. Note that the restriction is vacuous if one assumes bounded inpulse responses (in the sense of condition F-BIR in Pavan, Segal, and Toikka, 2014) and bounded output.

For all $t$, all $\theta^t \in \Theta^t$, all $\hat{\theta}_t \in \Theta_t$, let

$$D_t(\theta^t; \hat{\theta}_t) \equiv \mathbb{E}\left[\sum_{s\geq t} \delta^{s-t} I_{(t),s}(\theta^s) q_s(\theta^s_{-t}, \hat{\theta}_t) \mid \theta^t\right]. \tag{11}$$

**Theorem 1 (Pavan, Segal, and Toikka, 2014)** *Suppose the process $F$ is regular. The mechanism $\chi = \langle \mathbf{q}, \mathbf{p}\rangle$ is incentive compatible if and only if, (a) for all $t \geq 1$, all $\theta^{t-1}$, $V_t^A(\theta^t)$ is equi-Lipschitz continuous in $\theta_t$ with*

$$\frac{\partial V_t^A(\theta^t)}{\partial \theta_t} = \mathbb{E}\left[\sum_{s\geq t} \delta^{s-t} I_{(t),s}(\theta^s) q_s(\theta^s) \mid \theta^t\right] \text{ a.e. } \theta_t \in \Theta_t, \quad \text{(ICFOC)}$$

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

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

Note that the result in Theorem 1 is not specific to quasilinear settings. Conditions (ICFOC) and (Int-M) characterize the entire set of implementable allocations also in settings in which payments are either absent, or agents' payoffs are not linear in the payments (see also the discussion in Section 6, below). Also observe that, in quasilinear settings, there always exist payments $\mathbf{p}$ that guarantee that the (ICFOC) conditions are satisfied at all histories (see Pavan, Segal, and Toikka, 2014 for the construction of such payments, and Corollary 2, below, for a discussion of the extent to which such payments are unique). In such quasilinear environments, (ICFOC) imposes restrictions on the payments $\mathbf{p}$, while the integral monotonicity conditions (Int-M) impose restrictions on the output schedule $\mathbf{q}$.

**Sketch of the proof.** The formal arguments supporting the claim that (ICFOC) is a necessary condition for incentive compatibility are in Pavan, Segal, and Toikka (2014). For a heuristic intuition, see the discussion for the case $T = 2$ preceding the theorem. Here I focus on the necessity of the integral monotonicity conditions (Int-M) and on the fact that (ICFOC) and (Int-M), jointly, imply incentive compatibility. For this purpose, fix $t$ and $\theta^{t-1}$ and drop $\theta^{t-1}$ to ease the exposition. Let $U_t(\hat{\theta}_t; \theta_t)$ denote the agent's continuation payoff when his period-$t$ type is $\theta_t$, he reports $\hat{\theta}_t$ in period $t$, and then reports truthfully at all subsequent periods. Observe that $U_t(\theta_t; \theta_t) = V_t^A(\theta_t)$. Incentive compatibility then requires that , for all $\hat{\theta}_t \in \Theta_t$,

$$\hat{\theta}_t \in \arg\max_{\theta_t} \left\{U_t(\hat{\theta}_t; \theta_t) - V_t^A(\theta_t)\right\}.$$

Next, note that (ICFOC) implies that, for $\hat{\theta}_t$ fixed, the function $g_t(\cdot; \hat{\theta}_t) : \Theta_t \to \mathbb{R}$ given by $g_t(\theta_t; \hat{\theta}_t) \equiv U_t(\hat{\theta}_t; \theta_t) - V_t^A(\theta_t)$ is Lipschitz continuous with derivative equal to

$$g_t'(\theta_t; \hat{\theta}_t) = \partial U_t(\hat{\theta}_t; \theta_t)/\partial \theta_t - dV_t^A(\theta_t)/d\theta_t$$

for almost all $\theta_t$. If the mechanism $\chi = \langle \mathbf{q}, \mathbf{p} \rangle$ is incentive compatible, then reporting truthfully from period $t+1$ onwards must be optimal for the agent after reporting truthfully in period $t$. Because the environment is Markov (that is, payoffs separate over time and the process governing the evolution of the agent's type is Markov), then reporting truthfully from period $t+1$ onward must also be optimal after any deviation in period $t$. Consider then a fictitious environment in which the agent's period-$t$ report is exogenously fixed at $\hat{\theta}_t$. Incentive compatibility of the mechanism $\chi = \langle \mathbf{q}, \mathbf{p} \rangle$ in the primitive environment implies incentive compatibility of the same mechanism in the fictitious environment. The same arguments that establish the necessity of condition (ICFOC) in the primitive environment then also imply that, in the fictitious environment, $U_t(\hat{\theta}_t; \theta_t)$ must be equi-Lipschitz continuous in $\theta_t$ with $\partial U_t(\hat{\theta}_t; \theta_t)/\partial \theta_t = D_t((\theta^{t-1}, \theta_t); \hat{\theta}_t)$ for almost all $\theta_t$. Because $dV_t^A(\theta_t)/d\theta_t = D_t((\theta^{t-1}, \theta_t); \theta_t)$, one then has that

$$g_t(\hat{\theta}_t; \hat{\theta}_t) - g_t(\theta_t; \hat{\theta}_t) = \int_{\theta_t}^{\hat{\theta}_t} \frac{\partial g_t(x; \hat{\theta}_t)}{\partial \theta_t} dx = \int_{\theta_t}^{\hat{\theta}_t} \Big[ D_t((\theta^{t-1}, x); \hat{\theta}_t) - D_t((\theta^{t-1}, x); x) \Big] dx. \tag{12}$$

Hence, $\hat{\theta}_t \in \arg\max_{\theta_t} \left\{ g_t(\theta_t; \hat{\theta}_t) \right\}$ only if (Int-M) holds. This establishes that (ICFOC) and (Int-M) are jointly necessary. To see that they are also jointly sufficient, observe that, when these conditions are satisfied, no single one-stage deviation from truthful reporting is optimal (to see this, note that (ICFOC) and (Int-M) imply that $U_t(\hat{\theta}_t; \theta_t) \leq V_t^A(\theta_t)$ all $\theta_t, \hat{\theta}_t \in \Theta_t$ all $t$, all $\theta^{t-1}$). Along with the facts that (a) payoffs are continuous at infinity, and (b) the environment is Markov, the above result then implies that all other deviations are also suboptimal.[20] QED

In a Markov environment, the combination of integral monotonicity (Int-M) with the envelope representation of the equilibrium payoffs (ICFOC) thus fully characterizes incentive compatibility. That the environment is Markov implies that the agent's incentives in any period depend only on his current true type and his past reports, but not on his past true types. In turn this implies that, when a single departure from truthful reporting is suboptimal, then truthful reporting (at all histories) dominates any other strategy. Importantly, note that every environment can be "Markovized" by expanding the state space (e.g., by letting the state be equal to $\omega_t = \theta^t$ all $t$). In this case, Myerson's (1986) revelation principle takes the form of Doepke and Townsend's (2004) revelation principle – without loss of generality, one can restrict attention to mechanisms in which the agent reports $\omega_t = \theta^t$ at each period. When the agent is asked to report $\omega_t = \theta^t$ in all periods, it is always without loss of generality to consider

[20] Recall that the Markov assumption implies that, when no single deviation from truthful reporting is profitable at any truthful history (i.e., conditional on having reported truthfully in the past) then single deviations from truthful reporting are also suboptimal at all other non-truthful histories.

mechanisms in which the agent reports truthfully at all periods, irrespective of his past behavior. This property, however, does not necessarily make the characterization of incentive compatibility easier. The advantage is that it allows one to focus on "strongly truthful" strategies (that is, strategies prescribing truthful reporting at all histories); The disadvantage is that one needs to deal with incentive compatibility on a multidimensional state space. In certain environments, the dimensionality of the state space is irrelevant, so "Markovizing" the state space is the way to go (this is the case, for example, when implementing efficiency using VCG/AGV-type of mechanisms, or when maximizing profits in environments in which virtual surplus takes the form of an additive or multiplicative transformation of the true surplus, as in Kakade et al., 2013 or in Fershtman and Pavan, 2016). Working with this alternative state representation is also useful in environments in which the agent possesses private information not only about the realizations of a process but also about the process itself (which amounts to assuming that the agent's private information is multidimensional, with a component that is fully persistent – see, e.g., Boleslavsky and Said, 2013).

Static versions of the integral-monotonicity conditions appeared in the literature on implementability (see Rochet, 1987 or Carbajal and Ely, 2013 and the references therein). This condition generalizes the more familiar monotonicity conditions typically encountered in static settings with supermodular payoffs, unidimensional types, and unidimensional decisions by which an allocation rule is implementable if and only if it is monotone. As the above integral monotonicity condition reveals, what is required by incentive compatibility in more general settings is 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 NPV of expected future output, *discounted by impulse responses*

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

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 it is sufficiently monotone, on average, where the average is both over states and time.

The sufficiency results in the literature are typically based on stronger notions of monotonicity (see, for example, Battaglini, 2005, or Eső and Szentes, 2007) with the exception of the mean-preserving-spread case of Courty and Li (2000). However, there are interesting environments where the optimal allocation rule fails to be strongly monotone (that is, where the nonmonetary allocation in certain periods fails to be increasing in either the current or past reports), and/or where the kernels naturally depend on past decisions, or fail first-order stochastic dominance. For instance, the optimal allocation rule in Pavan, Segal, and Toikka (2014)'s bandit auction, as well as the optimal matching rule in Fershtman and Pavan (2016) fail to be strongly monotone and yet they satisfy integral monotonicity.

The above result has two important implications, which are summarized in the next two corollaries.

**Corollary 2 (payments equivalence)** *Let* $\chi = \langle \mathbf{q}, \mathbf{p} \rangle$ *and* $\bar{\chi} = \langle \mathbf{q}, \bar{\mathbf{p}} \rangle$ *be any two mechanisms implementing the same non-monetary decisions* $\mathbf{q}$. *There exists a constant* $K$ *such that for* $F$*-almost every* $\theta^T$

$$\sum_{t=1}^{T} \delta^t p_t(\theta^t) = \sum_{t=1}^{T} \delta^t \bar{p}_t(\theta^t) + K.$$

**Corollary 3 (irrelevance result)** *Let* $\tilde{\chi} = \left\langle \tilde{\mathbf{q}}, \tilde{\mathbf{p}} \right\rangle$ *be an incentive-compatible mechanism for the environment in which the agent is asked to report the orthogonal shocks* $\varepsilon$ *starting from* $t = 2$. *Suppose now the principal can observe the shocks* $\varepsilon$. *In any mechanism* $\tilde{\chi}' = \left\langle \tilde{\mathbf{q}}, \tilde{\mathbf{p}}' \right\rangle$ *implementing the same non-monetary decisions* $\tilde{\mathbf{q}}$, *there exists a constant* $K$ *such that, for any* $\theta_1 \in \Theta_1$,

$$\mathbb{E}\left[\sum_{t=1}^{T} \delta^t \tilde{p}_t(\theta_1, \varepsilon)\right] = \mathbb{E}\left[\sum_{t=1}^{T} \delta^t \tilde{p}'_t(\theta_1, \varepsilon)\right] + K.$$

The first result is a strong version of revenue-equivalence that says that any two mechanisms implementing the same non-monetary decisions yield the same revenues, not just in expectation, but in each state of the world. The result is obtained by using inductively the necessary envelope conditions in (ICFOC) at all histories. In other words, in each state of the world, the NPV of the payments is pinned down by the quantity schedule $\mathbf{q}$ and the payoff of the lowest period-1 type, $V_t^A(\underline{\theta}_1)$. This stronger version of payment-equivalence is particularly relevant for settings in which the utility the agent derives from the payments is not linear, as in most managerial compensation and taxation models (see, e.g., Farhi and Werning, 2013, Garrett and Pavan, 2015, and Makris and Pavan, 2016). For a discussion of how the above result extends to settings with multiple agents, see Pavan, Segal, and Toikka (2014).

The second result says that the observability of the shocks is irrelevant for the principal's payoff. The two results are obviously related in the sense that they both follow from the necessary conditions for incentive compatibility summarized by the envelope representation in (ICFOC). For a discussion of how the irrelevance result in the above corollary extends to certain environments combining adverse selection with moral hazard, see Eső and Szentes (2015). For a discussion of how this result may fail when the agents' initial type is drawn from a discrete distribution, see Pavan (2007) and Krähmer and Strausz (2015b).

### 4.2 Full and Relaxed Programs and the First-Order Approach

The results discussed above can be used to arrive at properties of optimal dynamic mechanisms. In particular, the complete characterization of the set of incentive compatible rules $\chi = \langle \mathbf{q}, \mathbf{p} \rangle$ in the previous theorem implies that the principal's *full program* can be stated as follows:

$$\mathcal{P}: \begin{cases} \max_{\chi=\langle \mathbf{q},\mathbf{p}\rangle} \mathbb{E}\left[\sum_t \delta^{t-1}(p_t(\theta^t) - C(q_t(\theta^t))\right] \\ \text{subject to} \\ V_1^A(\theta_1) \geq 0 \text{ all } \theta_1, \\ \text{ICFOC-}(t)\text{: } V_t^A(\theta^{t-1}, \theta_t) \text{ abs. cont. in } \theta_t \text{ with } \frac{\partial V_t^A(\theta^t)}{\partial \theta_t} = D_t(\theta^t; \theta_t) \text{ a.e. } \theta_t, \\ \quad \text{all } t, \text{ all } \theta^{t-1}, \\ \text{Int-M: } \int_{\hat{\theta}_t}^{\theta_t} [D_t((\theta^{t-1}, x); x) - D_t((\theta^{t-1}, x); \hat{\theta}_t)]dx \geq 0 \text{ all } t, \text{ all } (\theta^t, \hat{\theta}_t), \end{cases}$$

with the functions $D_t(\theta^t; \hat{\theta}_t)$ as defined in (11). Solving the above problem can still be tedious. The approach followed in most applied papers consists in solving a relaxation of the above program, in which all the envelope conditions ICFOC-($t$), for $t > 1$, all the integral-monotonicity conditions, and all the participation constraints but the one for the lowest period-1 type are dropped. Dropping all these constraints leads to the following *relaxed program:*

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

Once the solution to the relaxed program is at hand, by "reverse engineering," one then identifies primitive conditions (on the process $F$ and, in more general settings, on the utility functions) guaranteeing that the remaining constraints are also satisfied. This approach is often referred to as "first-order approach."

Using ICFOC-(1), and integrating by parts, one can rewrite the principal's objective as "*Dynamic Virtual Surplus*"

$$\mathbb{E}\left[\sum_t \delta^{t-1}\left(\left(\theta_t - \frac{1-F_1(\theta_1)}{f_1(\theta_1)} I_t(\theta^t)\right) q_t(\theta^t) - C(q_t(\theta^t))\right)\right] - V_1^A(\underline{\theta}). \quad \text{(DVS)}$$

Hence, in any incentive-compatible mechanism, irrespective of whether or not the mechanism solves the relaxed program, the principal's payoff takes the form of Dynamic Virtual Surplus. The latter combines the sum of the principal and the agent's gross payoffs with handicap terms that account for the cost to the principal of leaving information rents to the agents. In a dynamic setting, such handicaps combine the familiar term from static mechanism design, $[1 - F_1(\theta_1)]/f_1(\theta_1)$, with the impulse response functions $I_t(\theta^t)$ that link the period-$t$ type to the period-1 type. While the period-1 inverse hazard rate $[1 - F_1(\theta_1)]/f_1(\theta_1)$ controls for the importance the principal assigns to rent-extraction relative to efficiency, the impulse responses $I_t(\theta^t)$ control for the effect of distorting period-$t$ output on the agent's expected rent, as perceived at the moment of contracting, i.e., in period 1.

In the context of the simple screening problem under examination here, the Dynamic Virtual Surplus function can be maximized history by history. Assume, for example, that $C$ is quadratic and that $\mathcal{Q} = \mathbb{R}_+$, as in the discrete example in the previous section. The allocation rule that maximizes (DVS) is such that, for any $t$, any $\theta^t$

$$q_t(\theta^t) = \max\left\{\theta_t - \tfrac{1-F_1(\theta_1)}{f_1(\theta_1)} I_t(\theta^t); 0\right\}. \tag{13}$$

These optimality conditions are the continuous-type analogs of the conditions in the discrete-type example in the previous section. In particular, they imply (a) "no-distortion at the top" (that is, $q_t(\theta^t) = q_t^{FB}(\theta^t)$ if $\theta_1 = \bar{\theta}$, all $t$, all $\theta_{-1}^t$) and (b), in case 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 $I_t(\theta^t) \neq 0$. As in the example in the previous section, the dynamics of distortions are then driven by the dynamics of the impulse responses of future types to the initial ones. The smaller the impulse responses, the smaller the distortions. Furthermore, if impulse responses decline, on average, over time, so do the distortions.

The formula in (13) provides useful information about how the principal distorts the provision of output over time. However, the formula is valid only insofar as the solution to the relaxed program satisfies all the remaining constraints of the full program. It is easy to see that, when output is non-negative and the process satisfies first-order stochastic dominance, the solution to the relaxed program satisfies the remaining participation constraints for all period-1 types above the lowest one. In fact, ICFOC-(1) implies that

$$V_1^A(\theta_1) = V_1^A(\underline{\theta}) + \int_{\underline{\theta}}^{\theta_1} \mathbb{E}\left[\sum_{t\geq 1} \delta^{t-1} I_t(\theta^t) q_t(\theta^t) \mid x\right] dx \geq V_1^A(\underline{\theta})$$

for all $\theta_1 > \underline{\theta}$. Also, note that the remaining ICFOC-($t$) constraints for $t > 1$ neglected in the relaxed program can always be satisfied by letting the payments be equal to

$$p_t\left(\theta^t\right) = \theta_t q_t(\theta^t) - \int_{\underline{\theta}}^{\theta_t} \mathbb{E}\left[\sum_{s\geq t} \delta^{s-t} I_{(t),s}(\theta^s) q_s(\theta^s) \mid (\theta^{t-1}, x)\right] dx$$
$$+ \delta \mathbb{E}\left[\int_{\underline{\theta}}^{\theta_{t+1}} \mathbb{E}\left[\sum_{s\geq t+1} \delta^{s-t-1} I_{(t+1),s}(\theta^s) q_s(\theta^s) \mid (\theta^t, x)\right] dx \mid \theta^t\right].$$

Finally, note that the remaining (Int-M) constraints are satisfied if the quantity schedule $\mathbf{q}$ that solves the above relaxed program is sufficiently monotone, in the sense that

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

is nonincreasing in $\hat{\theta}_t$ all $t$ all $\theta^t$. Using (13), one can easily verify that, when the cost is quadratic and output is nonnegative, these conditions are trivially satisfied if, for example, the agent's type follows an AR(1) process (in which case, $I_t(\theta^t) = \gamma^{t-1}$) and the period-1 distribution $F_1$ is log-concave. Interestingly, note that, in this case, the volatility of the shocks $\varepsilon_t$ is irrelevant for the dynamics of the distortions. I refer the reader to Pavan, Segal, and Toikka (2014) for other examples where the solution to the relaxed program solves the full program and for a more detailed discussion of how the dynamics of the impulse responses drive the dynamics of the distortions when the output schedules in the optimal contracts are the ones that solve the relaxed program.

## 5 ROBUST PREDICTIONS: VARIATIONAL APPROACH

As mentioned in the Introduction, the predictions about the dynamics of distortions under optimal contracts identified in the literature are confined to environments in which the output schedules coincide with those that solve the relaxed program (equivalently, to settings in which the "first order approach" is valid). In this section, I illustrate an alternative approach, which does not permit one to fully characterize the optimal schedules, but offers a way of identifying certain predictions for the dynamics of the average distortions that do not rely on technical assumptions on the stochastic process governing the evolution of the agents' private information or the cost function made to validate the first-order approach. The first paper to follow such an approach is Garrett and Pavan (2015), in the context of managerial compensation. That paper is discussed in Section 6, below. The exposition in this section follows from Garrett, Pavan, and Toikka (2016), where the environment is similar to the one under consideration here.

For simplicity, assume $\Theta_t = \Theta = [\underline{\theta}, \bar{\theta}]$, for all $t$, and then denote by $\mathcal{B}(\Theta)$ the Borel sigma-algebra over $\Theta$. For any $A \in \mathcal{B}(\Theta)$, any $t \geq 1$, then let $P^t(A;\theta)$ denote the probability that $\theta_t \in A$ given that $\theta_1 = \theta$.

**Definition 2** *The type process $F$ is* ergodic *if there exists a unique (invariant) probability measure $\pi$ on $\mathcal{B}(\Theta)$ whose support has a nonempty interior such that for all $\theta \in \Theta$,*

$$\sup_{A\in\mathcal{B}(\Theta)}\left|P^t(A;\theta)-\pi(A)\right| \to 0 \textit{ as } t\to\infty. \tag{14}$$

**Definition 3** *The type process $F$ is* stochastically monotone *if, for all $t \geq 2$, $\theta' > \theta \Rightarrow F_t(\cdot \mid \theta') \succsim_{FOSD} F_t(\cdot \mid \theta)$.*

**Definition 4** *The type process $F$ is* stationary *if $F_1 = \pi$ and $F_t = F_s$, all $t, s > 1$.*

The above properties are "economic properties" as opposed to technical conditions meant to validate certain first-order conditions. For example, ergodicity captures the idea that the probability the agent assigns to his type reaching a certain level in the distant future is invariant in his current type. Likewise, the property that $F$ is stochastically monotone captures the idea that an agent who, in the present period, assigns a higher marginal value to the principal's product than another agent, expects to derive a higher value also in the next period. Finally, stationarity captures the idea that, at the time of contracting, the process has already evolved for long enough to have converged to the invariant distribution. The nature of these conditions is very different from, say, the log-concavity of the period-1 distribution $F_1$, or the monotonicity of the impulse response functions $I_t(\theta^t)$ in the realized types, which are technical conditions assumed in applications to validate the first-order approach (via reverse engineering), but which do not have a strong economic appeal.

The aim here is to identify predictions about the dynamics of distortions under optimal contracts that do not hinge on the validity of the first-order approach. In general, solving the full program is hard, if not impossible. The idea in Garrett, Pavan, and Toikka (2016) is that various properties of optimal contracts can be identified without fully solving for the optimal contracts. To see this, suppose that $\chi = \langle \mathbf{q}, \mathbf{p} \rangle$ is an optimal mechanism. Then use the results in the previous section to observe that any perturbation $\mathbf{q}'$ of the output schedule $\mathbf{q}$ that preserves (Int-M) along with any adjustment $\mathbf{p}'$ of the payment schedule $\mathbf{p}$ that preserves (ICFOC) and the period-1 participation constraints yield a mechanism $\chi' = \langle \mathbf{q}', \mathbf{p}' \rangle$ that is individually rational and incentive compatible. For the original mechanism $\chi = \langle \mathbf{q}, \mathbf{p} \rangle$ to be optimal, it must be that any such perturbation is unprofitable. One can then use such perturbations to identify properties of the optimal mechanisms.[21]

To illustrate, assume the period-1 participation constraints bind only at the bottom, i.e., only for $\theta_1 = \underline{\theta}$ (note that this is always the case when the process is stochastically monotone and output is nonnegative). For a moment, assume also interior solutions. A simple perturbation that preserves (Int-M) is to add a constant $a \in \mathbb{R}$ to the entire period-$t$ output schedule $q_t(\theta^t)$. Then, use the fact that, under any incentive-compatible mechanism, the principal's payoff must coincide with Dynamic Virtual Surplus (DVS), to verify that a necessary condition for the optimality of the output schedule $\mathbf{q}$ is that the derivative of Dynamic Virtual Surplus (DVS) with respect to $a$ evaluated at $a = 0$ must vanish. In the context of the economy under consideration here, the above requirement is fulfilled if and only if

$$\mathbb{E}\left[\theta_t - \tfrac{1-F_1(\theta_1)}{f_1(\theta_1)} I_t(\theta^t)\right] = \mathbb{E}\left[C'(q_t(\theta^t))\right]. \tag{15}$$

[21] The existence of optimal mechanisms is proved in Garrett, Pavan, and Toikka (2016).

The left-hand side in (15) is the *average* marginal benefit of expanding the period-$t$ output uniformly across the period-$t$ histories. The benefit is in virtual terms to account for the effect of higher output on the surplus the principal must leave to the agent at the time of contracting (recall that the impulse responses $I_t = I_{(1),t}$ are the ones with respect to the period-1 types). The right-hand side is the *average* marginal cost. Optimality requires that the average period-$t$ distortion be equal to the average period-$t$ handicap, where the former is given by

$$\mathbb{E}[\text{period-}t \text{ distortion}] \equiv \mathbb{E}[\theta_t - C'(q_t(\theta^t))]$$

whereas the latter is given by

$$\mathbb{E}[\text{period-}t \text{ handicap}] \equiv \mathbb{E}\left[\frac{1-F_1(\theta_1)}{f_1(\theta_1)} I_t(\theta^t)\right].$$

The optimality condition in (15) is thus the analog of the optimality condition (13) describing the solution to the relaxed program, but with the condition required to hold only *in expectation* as opposed to pathwise.

To derive predictions about the dynamics of the average distortions, one then investigates the dynamics of the average handicap.

**Theorem 2 (Garrett, Pavan, and Toikka, 2016)** *Assume the process $F$ is regular and ergodic. Then*

$$\mathbb{E}\left[\frac{1-F_1(\theta_1)}{f_1(\theta_1)} I_t(\theta^t)\right] \to 0 \text{ as } t \to \infty.$$

*Moreover, if $F$ is stochastically monotone, convergence is from above. If, in addition, $F$ is stationary, then convergence is monotone in time.*

Hence, the dynamics of the average handicaps are entirely driven by three economic properties of the type process: ergodicity, stationarity, and stochastic monotonicity. When output is interior and participation constraints bind only at the bottom, the above result thus implies that, when the process is ergodic, average distortions vanish in the long run. If, in addition, the process is stochastically monotone, average distortions vanish from above, meaning that the average distortion is positive (equivalently, output is distorted downward relative to the first best). Finally, if, in addition, the initial distribution coincides with the invariant distribution, then convergence is monotone in time, meaning that the average period-$t$ distortion is larger than the average period-$s$ distortion, for any $s > t$. The above predictions can be generalized as follows.

**Definition 5** *Given the output process* $\mathbf{q}$, *period-$t$ output is* strictly interior *if there exists $\varepsilon_t > 0$ such that*

$$\min \mathcal{Q} + \varepsilon_t \le q_t(\theta^t) \le \max \mathcal{Q} - \varepsilon_t \quad \textit{for all } \theta^t.$$

*The output schedule* $\mathbf{q}$ *is* strictly interior *if the above condition holds for all $t$. It is* eventually strictly interior *if the condition holds for all sufficiently large $t$.*

**Definition 6** *Given the output process* $\mathbf{q}$*, the period-t distortions are* downwards *if* $\theta_t \geq C'(q_t)$ *almost surely. Distortions are* eventually downward *if the above condition holds for all sufficiently large* $t$.

**Theorem 3 (Garrett, Pavan, and Toikka, 2016)** *Suppose* $F$ *is regular and ergodic. (a) If* $\mathbf{q}$ *is eventually strictly interior, then* $\lim_{t\to\infty} \mathbb{E}[\theta_t - C'(q_t)] = 0$. *(b) If distortions are eventually downward, then* $q_t \xrightarrow{p} q_t^{FB}$.

Hence, if output is eventually strictly interior (which is necessarily the case when $\mathcal{Q} = \mathbb{R}$, that is, when the principal can either buy or sell output to the agent, as in certain trading models), then distortions eventually vanish in expectation. When, instead, $\mathcal{Q} = \mathbb{R}_+$, that is, when the principal can sell to, but not buy from the agent, then

$$\limsup_{t\to\infty} \mathbb{E}[\theta_t - C'(q_t)] \leq 0.$$

In this case, long-run average distortions are either zero or upward. Together, the above results imply that, if convergence to the first best fails, it must be because, eventually, within the same period, certain types over-consume, while others are completely excluded, and this pattern must occur infinitely often.

Part (b) in Theorem 3 in turn implies that, when distortions are eventually downward, as is always the case when $F$ is strongly monotone and the optimal schedules solve the relaxed program, then convergence to the first-best is not just in expectation but in probability.

Importantly, note that the variational approach briefly described in this section (and which is still in its infancy) is different from the one in the earlier literature, and which is extensively used in the new dynamic public finance literature. The earlier approach perturbs the agent's compensation over multiple periods, holding fixed the agent's expected payoff at any history (thus trivially preserving incentive compatibility). Such an approach yields an inverse Euler equation that links average compensation over consecutive periods (see, e.g., Rogerson, 1985). It does not permit one, however, to identify predictions about the dynamics of the non-monetary allocations (e.g., output or effort). The new approach, instead, perturbs the agent's continuation payoff across different type histories. As such, the new approach builds on Theorem 1, above, to identify perturbations of the non-monetary allocations that preserve incentive compatibility.

## 6 BEYOND THE QUASI-LINEAR CASE

The results reviewed above are for environments in which the utility the players derive from the numeraire is linear. In such environments, the way the principal distributes the payments over time is irrelevant. For certain applications, this is clearly not a desirable assumption. For example, both in the managerial compensation literature as well as in the optimal taxation literature it is customary to assume that the agent's utility over the payments he receives from the principal, or over his consumption, is nonlinear. In this section, I show how the above results must be adapted to accommodate for non-quasi-linear payoffs. In the context of the environment examined thus far, such an extension could be accommodated by letting the agent's payoff be equal to $U^A \equiv \sum_{t=1}^{T} \delta^{t-1} \left[\theta_t q_t - v^A(p_t)\right]$ for some concave function $v^A(\cdot)$. It is easy to see that this enrichment does not affect the characterization of incentive compatibility, which continues to be determined by Conditions (ICFOC) and (Int-M) in Theorem 1. The point where the analysis departs from the discussion above is in the characterization of the optimal policies. When one of the two players' payoff is non-quasi-linear (that is, in the absence of transferable utility), the optimal policies cannot be obtained via simple point-wise maximization of the Dynamic Virtual Surplus function. Furthermore, the first-best allocations cannot be described in terms of the usual equalization of marginal cost of production to marginal value, for one must take into account the non-transferability of the principal's cost to the agent (or, equivalently, of the agent's utility from consuming the good to the principal).

In this section, I show how the Myersonian and the variational approaches discussed in the previous sections must be adapted to accommodate for these complications. To facilitate the connection to the pertinent literature, I consider a slightly different environment that is meant to capture, in reduced form, the non-quasi-linear applications typically considered in the literature. To this purpose, suppose that the principal's and the agent's payoffs are now given by

$$U^P \equiv \sum_{t=1}^{T} \delta^{t-1} \left[y_t - c_t\right] \text{ and } U^A \equiv \sum_{t=1}^{T} \delta^{t-1} \left[v(c_t) - \psi(y_t, \theta_t)\right].$$

In the case of managerial compensation, one can interpret $y_t$ as the cash flows the manager generates for the firm, $c_t$ as the agent's period-$t$ consumption, and $\psi(y_t, \theta_t)$ as the disutility the agent derives from generating cash flows $y_t$ when his period-$t$ productivity is $\theta_t$. In a procurement/regulation model, $y_t$ can be interpreted as the gross surplus the regulated firm generates to society, $c_t$ as the period-$t$ compensation paid by the regulator to the firm, and $\psi(y_t, \theta_t)$ as the cost incurred by the firm when its period-$t$ efficiency parameter is $\theta_t$. In the case of nonlinear income taxation, one can interpret $y_t$ as the income produced by the worker in period $t$, $c_t$ as the period-$t$ consumption (i.e., the net-of-tax disposable income), $y_t - c_t$ as the tax collected by the government, and $\theta_t$ as the worker's effective wage (or productivity). In the first two applications, the principal's problem consists in maximizing the ex-ante expectation of $U^P$ subject to the usual incentive compatibility and participation constraints. In the case of optimal taxation, the planner's dual problem consists in maximizing tax revenues subject to the usual IC constraints and the constraint that

$$\int q(V^A(\theta_1))dF_1(\theta_1) \geq \kappa,$$

for some $\kappa \geq 0$, where the function $q(\cdot)$ captures the nonlinear Pareto weights that the planner assigns to different period-1 types (see also the discussion in the next section).

Here, I focus on the case of managerial compensation. As in Garrett and Pavan (2015), assume that

$$\psi(y_t, \theta_t) = \frac{1}{2}(y_t - \theta_t)^2 \text{ and } \theta_t = \gamma\theta_{t-1} + \varepsilon_t$$

with $\varepsilon_t$ drawn from an absolutely continuous distribution $G_t$ with mean zero, independently across $t$. The difference $e_t = y_t - \theta_t$ should be interpreted as the agent's period-$t$ effort. For simplicity, assume $T = 2$. Both $(\theta_1, \theta_2) \in \Theta_1 \times \Theta_2$ and $(e_1, e_2) \in \mathbb{R}^2$ are the manager's private information. Instead, the cash flows $y \equiv (y_1, y_2)$ are verifiable, and hence can be used as a basis for the manager's compensation. The function $v : \mathbb{R} \to \mathbb{R}$ is strictly increasing, weakly concave, surjective, Lipschitz continuous, and differentiable. The case where $v$ is linear corresponds to the case where the manager is risk-neutral, while the case where $v$ is strictly concave corresponds to the case where he is strictly risk-averse. Let $w \equiv v^{-1}$ and $\psi_y(y_t, \theta_t) \equiv \partial\psi(y_t, \theta_t)/\partial y_t$.

That productivity follows an autoregressive process (in which case the impulse responses of period-2 types, $\theta_2$, to period-1 types, $\theta_1$, are constant and equal to $\gamma$) simplifies the exposition, but is not essential for the results. The perturbations considered in Garrett and Pavan (2015) preserve incentive compatibility under more general processes in which impulse responses are neither constant nor monotone. I will refer to $\gamma = 1$ as to the case of "full persistence" (meaning that, holding effort fixed, the effect of any shock to period-1 productivity on the firm's average cash flows is constant over time). To ease the comparison to the results in the previous section, I describe a mechanism $\chi = \langle \mathbf{y}, \mathbf{c} \rangle$ in terms of its cash flows and compensation policy. I refer the reader to Garrett and Pavan (2015) for a discussion of how one can alternatively describe a mechanism in terms of recommended effort policy and a payment scheme (with the latter defined also for off-path cash flow observations).

The result below parallels the one in Theorem 1, above, for the case of transferable utility by providing a complete characterization of incentive compatibility. Let $\theta \equiv (\theta_1, \theta_2)$ and for any $(\theta; \mathbf{y})$, let

$$\begin{aligned} W(\theta; \mathbf{y}) \equiv\ & \psi(y_1(\theta_1), \theta_1) + \psi(y_2(\theta), \theta) \\ & + \int_{\underline{\theta}_1}^{\theta_1} \left\{ \psi_y(y_1(s), s) + \gamma \mathbb{E}\left[\psi_y(y_2(s, \theta_2), \theta_2) \mid s\right] \right\} ds \\ & + \int_{\underline{\theta}_2}^{\theta_2} \psi_y(y_2(\theta_1, s), s) ds - \mathbb{E}\left[ \int_{\underline{\theta}_2}^{\theta_2} \psi_y(y_2(\theta_1, s), s) ds \mid \theta_1 \right]. \end{aligned} \tag{16}$$

**Theorem 4 (Garrett and Pavan, 2015)** *The mechanism $\chi = \langle \mathbf{y}, \mathbf{c} \rangle$ satisfies all incentive-compatibility and participation constraints if and only if the following conditions jointly hold: (A) for all $\theta \in \Theta$,*

$$v\left(c_1\left(\theta_1\right)\right) + v\left(c_2\left(\theta\right)\right) = W\left(\theta;\mathbf{y}\right) + K \tag{17}$$

*where* $K \geq 0$ *is such that*

$$\mathbb{E}\left[W\left(\theta;\mathbf{y}\right) - \psi(y_1(\theta),\theta_1) + \psi(y_2(\theta),\theta) \mid \theta_1\right] + K \geq 0 \quad \forall \theta_1 \in \Theta_1; \tag{18}$$

*and (B)(i) for all* $\theta_1, \hat{\theta}_1 \in \Theta_1$,

$$\int_{\hat{\theta}_1}^{\theta_1} \left\{\psi_y\left(y_1(\hat{\theta}_1), s\right) + \gamma\mathbb{E}\left[\psi_y\left(y_2(\hat{\theta}_1,\theta_2),\theta_2\right) \mid s\right]\right\} ds$$
$$\leq \int_{\hat{\theta}_1}^{\theta_1} \left\{\psi_y\left(y_1\left(s\right), s\right) + \gamma\mathbb{E}\left[\psi_y\left(y_2(s,\theta_2),\theta_2\right) \mid s\right]\right\} ds, \tag{19}$$

*and B(ii)* $y_1\left(\theta_1\right) + \gamma\mathbb{E}\left[y_2(\theta_1,\theta_2) \mid \theta_1\right]$ *is nondecreasing in* $\theta_1$ *and, for all* $\theta_1 \in \Theta_1$, $y_2\left(\theta_1,\theta_2\right)$ *is nondecreasing in* $\theta_2$.

Condition (A) is the strong form of payoff-equivalence alluded to in Corollary 1, above: in each state $\theta$, the utility the agent derives from the compensation he receives from the principal is uniquely pinned down by the policy $\mathbf{y}$, up to a scalar $K \geq 0$ chosen to guarantee participation. Such compensation must provide the agent a payoff, net of the disutility of effort,

$$v\left(c_1\left(\theta_1\right)\right) + v\left(c_2\left(\theta\right)\right) - \psi(y_1(\theta_1),\theta_1) - \psi\left(y_2(\theta),\theta_2\right)$$

equal to his period-1 expected payoff

$$V^A(\theta_1) = V^A(\underline{\theta}_1) + \int_{\underline{\theta}_1}^{\theta_1} \left\{\psi_y\left(y_1(s), s\right) + \gamma\mathbb{E}\left[\psi_y(y_2(s,\theta_2),\theta_2) \mid s\right]\right\} ds \tag{20}$$

augmented by a term

$$\int_{\underline{\theta}_2}^{\theta_2} \psi_y(y_2(\theta_1,s),s)ds - \mathbb{E}\left[\int_{\underline{\theta}_2}^{\theta_2} \psi_y(y_2(\theta_1,s),s)ds \mid \theta_1\right],$$

that guarantees that the manager has the incentives to report truthfully not only in period 1 but also in period 2. Note that the expectation of this last term vanishes when computed based on period-1's private information, $\theta_1$. Part (A) thus imposes a restriction on the compensation scheme and is the analog of condition (ICFOC) in Theorem 1, above, applied jointly to $t = 1$ and $t = 2$. Note that the surplus that type $\theta_1$ expects over and above the one expected by the lowest period-1 type $\underline{\theta}_1$ is increasing in the effort that the firm asks of managers with initial productivities $\theta_1' \in \left(\underline{\theta}_1, \theta_1\right)$ in each of the two periods. This surplus is necessary to dissuade type $\theta_1$ from mimicking the behavior of these lower types. Also note that the scalar $K$ in (17) corresponds to the expected payoff $V^A(\underline{\theta}_1)$ of the lowest period-1 type. Using (20), it is easy to see that, when the effort requested is nonnegative, then, if $\underline{\theta}_1$ finds it optimal to accept the contract, so does any manager whose initial productivity is higher. This property, however, need not hold in case the firm requests a negative effort from a positive-measure set of types.

Next consider condition (B). This condition is the analog of the integral monotonicity Condition (Int-M), above, applied to the environment under examination. It imposes restrictions on the cash flow policy that are independent of the manager's felicity function, $v$. In particular, Condition (B)(ii) combines the familiar monotonicity constraint for the second-period cash flows from static mechanism design (e.g., Laffont and Tirole, 1986) with a novel monotonicity constraint that requires the NPV of the expected cash flows, weighted by the impulse responses (which here are equal to one in the first period and $\gamma$ in the second period) to be nondecreasing in period-1 productivity.

Garrett (2015) show how the above results can be used to derive implications for the dynamics of distortions under optimal contracts. Because the environment features non-transferable utility, distortions are captured by *wedges*, as in the public finance literature.

**Definition 7 (wedges)** *For each* $t = 1, 2$ *and each* $\theta = (\theta_1, \theta_2)$*, the (local ex-post)* distortions *in the provision of incentives under an (incentive-compatible) mechanism* $\chi = \langle \mathbf{y}, \mathbf{c} \rangle$ *are given by the wedge*

$$D_t(\theta) \equiv 1 - \frac{\psi_y\left(y_t\left(\theta\right), \theta_t\right)}{v'\left(c_t\left(\theta\right)\right)} \tag{21}$$

*between the marginal effect of higher effort on the firm's cash flows and its marginal effect on the compensation necessary to preserve the manager's utility constant.*

As discussed above, the approach typically followed in the literature to identify optimal policies is the following. First, consider a *relaxed program* that replaces all incentive-compatibility constraints with Condition (17) and all individual-rationality constraints with the constraint that $K = V^A(\underline{\theta}_1) \geq 0$. Second, choose policies $(y_1, y_2, c_1)$ along with a scalar $K$ so as to maximize the firm's profits subject to the aforementioned constraints. However, recall that Condition (17) alone is necessary but not sufficient for incentive compatibility. Furthermore, when the solution to the relaxed program yields policies prescribing a negative effort over a positive-measure set of types, satisfaction of the participation constraint for the period-1 lowest type $\underline{\theta}_1$ does not guarantee satisfaction of all other participation constraints. Therefore, one must typically identify auxiliary assumptions on the primitives of the problem guaranteeing that the policies that solve the relaxed program are implementable.

The approach in Garrett and Pavan (2015) consists, instead, in using variational techniques analogous to those used in the previous section to arrive at properties of the optimal contracts, without fully solving for the optimal contracts. First, one uses the result in the previous theorem to express the principal's profits as

$$\mathbb{E}[U^P] = \mathbb{E}\left[y_1(\theta) + y_2(\theta) - c_1(\theta_1) - v^{-1}\left(W(\theta;\mathbf{y}) + K - v(c_1(\theta_1))\right)\right]. \tag{22}$$

One then proceeds by identifying "admissible perturbations" to implementable policies such that the perturbed policies remain implementable (i.e., continue to satisfy the conditions for incentive compatibility, as in the previous theorem). Finally, one uses such perturbations as test functions to identify properties of optimal policies.[22]

Through the above variational approach, Garrett and Pavan (2015) establish a series of results relating the dynamics of the expected wedges to (a) the agent's risk aversion, and (b) the persistence of the productivity process.

**Theorem 5 (Garrett and Pavan, 2015)** *Suppose the densities of the type process $F$ are bounded.*[23] *Fix the level of persistence $\gamma < 1$ of the manager's productivity, and assume that the manager's preferences over consumption in each period are represented by the function $v_\rho(\cdot)$, with lower values of $\rho$ indexing lower degrees of risk aversion.*[24] *Then there exists $\bar{\rho} > 0$ such that, for any $\rho \in [0, \bar{\rho}]$, under any optimal contract $|\mathbb{E}[D_2(\theta)]| < |\mathbb{E}[D_1(\theta_1)]|$, with $sign(\mathbb{E}[D_2(\theta)]) = sign(\mathbb{E}[D_1(\theta_1)])$.*

The result says that, for small degrees of risk aversion, average distortions decline over time, when the agent's type is less than fully persistent. While the result in the proposition focuses on the dynamics of distortions, the same properties apply to expected effort. Precisely, expected effort increases over time. Importantly, note that the result is established without imposing restrictions on the utility function $v$ and the period-1 distribution $F_1$ necessary to validate the first-order approach.

The intuition for why, for low degrees of risk aversion, distortions decline, on average, over time is best illustrated in the case of positive effort (for the case where effort takes on negative values, see Garrett and Pavan, 2015). Asking the manager to exert higher effort requires increasing the surplus that the firm must leave to the manager to induce him to reveal his productivity (this surplus is over and above the minimal compensation required to compensate

[22] For certain results, the perturbations in Garrett and Pavan (2015) consist in adding nonnegative constant functions $\alpha(\theta_1) = a > 0$ and $\beta(\theta) = b > 0$, all $\theta$, to the original cash flow policies $y_1(\theta_1)$ and $y_2(\theta)$ and then adjusting the compensation policy $c$ so that the payments continue to satisfy the conditions in the previous theorem. For other results, instead, they consider perturbations that preserve not only incentive compatibility but also the agent's expected payoff conditional on his period-1 type $\theta_1$. This is obtained by considering joint perturbations of $y_1$ and $y_2$ of opposite sign.

[23] Formally, there exist $a, b \in \mathbb{R}_{++}$ such that, for almost all $\theta_1 \in \Theta_1$, $\theta_2 \in \Theta_2(\theta_1)$, $a < f_1(\theta_1), f_2(\theta_2|\theta_1) < b$, where $\Theta_2(\theta_1) = Supp[F_2(\cdot|\theta_1)]$.

[24] Formally, let $(v_\rho)_{\rho \geq 0}$ be a collection of functions $v_\rho : \mathbb{R} \to \mathbb{R}$ with the following properties: (i) for each $\rho > 0$, $v_\rho$ is surjective, continuously differentiable, increasing, and strictly concave, with $v_\rho(0) = 0$ and $v'_\rho(0) = 1$; (ii) $v_0$ is the identity function; (iii) $v'_\rho$ converges to one, uniformly over $c$ as $\rho \to 0$.

the manager for his disutility of effort). In this case, the firm distorts downward the effort asked of those managers whose initial productivity is low to reduce the rents it must leave to those managers whose initial productivity is high. When productivity is not fully persistent, these distortions are more effective in reducing managerial rents when introduced in earlier periods than in later ones. Distortions are therefore smaller at later dates. The correction is most pronounced when productivity is least persistent. Indeed, as one approaches the case where productivity is independent over time (i.e., when $\gamma$ is close to zero), the expected effort the firm asks of each manager in the second period is close to the first-best level (here $e^{FB} = 1$).

The levels of risk aversion for which the result in the previous proposition holds (i.e., how large one can take $\bar{\rho}$) depends on the persistence $\gamma$ of the initial productivity $\theta_1$. For a fixed level of risk aversion, if $\gamma$ is close to 1, i.e., if the agent's productivity is highly persistent, then dynamics opposite to those in the previous theorem obtain: distortions increase, on average, over time, as illustrated in the next theorem.

**Theorem 6 (Garrett and Pavan, 2015)** *Assume the agent is strictly risk-averse.*

*(a) Suppose productivity is fully persistent (i.e., $\gamma = 1$). Then distortions weakly increase over time, i.e., for almost all $\theta_1$,*

$$\mathbb{E}\left[D_2(\theta) \mid \theta_1\right] \geq D_1(\theta_1). \tag{23}$$

*(b) Suppose densities and cash flows are uniformly bounded*[25] *and that, for $\gamma = 1$, the inequality in (23) is strict. There exists $\bar{\gamma} < 1$ such that, for all $\gamma \in \left[\bar{\gamma}, 1\right]$, (23) holds as a strict inequality.*

Again, to ease the discussion, suppose that the effort asked by the firm in each period is strictly positive and that distortions are nonnegative (the result in the proposition also applies to the case where the effort asked of certain types as well as the distortions are negative). When the manager is risk-averse, incentivizing high effort in period 2 is more costly for the firm than incentivizing the same effort in period one. This is because high effort requires a high sensitivity of pay to performance. When done in period 2, such a sensitivity exposes the manager to volatile compensation as a result of the manager's own private uncertainty about his period-2 productivity. Since the manager dislikes this volatility, he must be provided additional compensation by the firm. To save on managerial compensation, the firm then distorts, on average, period-2 incentives more than in period 1. Note that, when the agent's type is fully persistent (i.e., $\gamma = 1$), distortions increase over time, for any strictly concave $v$. The result extends to persistence levels that are sufficiently high.

[25] Formally, suppose there exists $b \in \mathbb{R}_{++}$ such that, for all $\theta_1 \in \Theta_1$, $\theta_2 \in \Theta_2(\theta_1)$, $f_2\left(\theta_2|\theta_1\right) < b$. In addition, suppose there exists $M \in \mathbb{R}_{++}$ and $\gamma' < 1$ such that, for all $\gamma \in [\gamma', 1]$, the optimal cash flow policy $y_t(\theta)$ is uniformly bounded (in absolute value) by $M$.

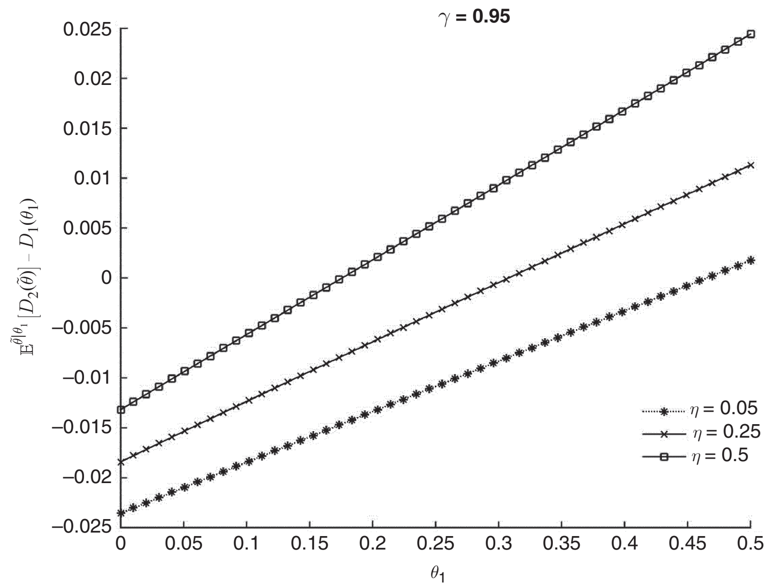


Figure 2 Differential between period-2 and period-1 distortions: $\gamma = 0.95$, $\eta = 0.05$, $\eta = 0.25$ and $\eta = 0.5$.

The above two results are illustrated in the two figures for the special case in which the optimal policies coincide with those that solve the relaxed program and the agent's utility function is isoelastic (i.e., $v(c) = (c^{1-\eta} - 1)/(1 - \eta)$). Figure 2 fixes the level of persistence to $\gamma = .95$ and shows how the difference $\mathbb{E}^{\tilde{\theta}|\theta_1}\left[D_2(\tilde{\theta})\right] - D_1(\theta_1)$ between the period-2 distortions and the period-1 distortions is affected by the degree of managerial risk aversion, $\eta$. As indicated in the two theorems above, higher degrees of risk aversion imply a higher differential between period-2 distortions and period-1 distortions. In particular, when $\eta = .05$ (that is, when the manager is close to being risk-neutral) the expected period-2 distortions are smaller than the period-1 distortions, for all but the period-1 very highest types. For higher degrees of risk aversion, instead, expected period-2 distortions are smaller than period-1 distortions over a smaller set of period-1 types $\theta_1$.

Figure 3, instead, fixes the coefficient of relative risk aversion to $\eta = 1/2$ and depicts the difference $\mathbb{E}^{\tilde{\theta}|\theta_1}\left[D_2(\tilde{\theta})\right] - D_1(\theta_1)$ between period-2 distortions and period-1 distortions for three different levels of persistence, $\gamma = .9$, $\gamma = .95$, and $\gamma = 1$. When productivity is perfectly persistent ($\gamma = 1$), the difference is strictly positive for all $\theta_1$. When, instead, $\gamma = .95$, or $\gamma = .9$, the difference continues to be positive, but only for sufficiently high values of $\theta_1$. That, for low values of $\theta_1$, the difference is negative reflects the fact that, for these types, period-1 effort is small. The firm can then afford to ask of these types a higher period-2 effort without having to pay them significant additional compensation.

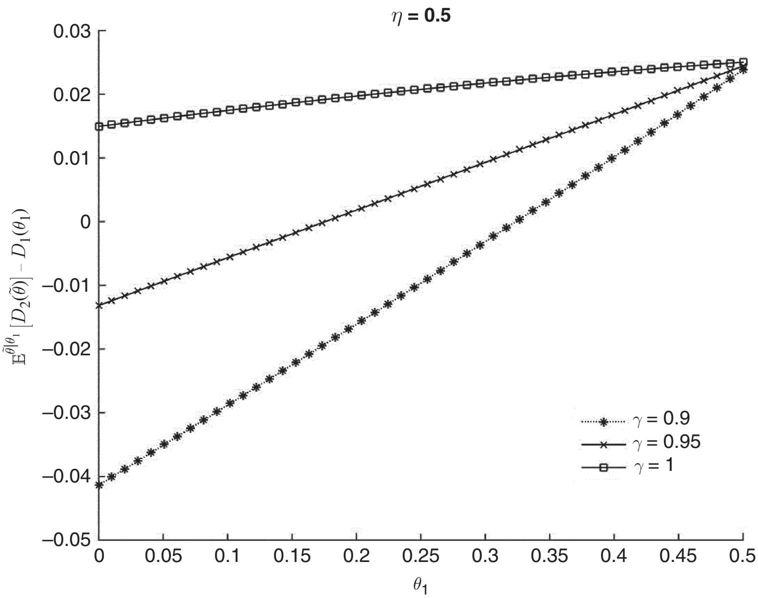


Figure 3 Differential between period-2 and period-1 distortions: $\eta = 1/2$, $\gamma = 0.9$, $\gamma = 0.95$ and $\gamma = 1$.

## 7 ENDOGENOUS TYPES

Most of the dynamic mechanism design literature assumes that the process governing the evolution of the agents' private information is exogenous. For many problems of interest this is a reasonable assumption. However, there are important applications in which accommodating for an endogenous process is essential. Consider, for example, the sale of experience goods. The buyers' valuations evolve over time as the result of the buyers' experimentation. The dynamics of the allocation of the good then determine the evolution of the buyers' valuations. Similar problems emerge in the presence of habit formation, addiction, and learning-by-doing.

The general model of Pavan, Segal, and Toikka (2014) is flexible enough to accommodate for endogenous (that is, decision-controlled) processes. That paper also offers an application to the design of dynamic auctions for the sale of experience goods. The authors show that the profit-maximizing mechanism takes the form of a sequence of bandit auctions where in each period the seller allocates the good (in limited capacity) to the bidders whose virtual Gittins index is the highest. The virtual index is the same as in the multi-arm bandit literature, but adjusted by a handicap that controls for the cost of informational rents. Interestingly, the optimal allocation rule is not strongly monotone, as a higher report at present, by inducing the agent to consume, may trigger a reduction in consumption in future periods in case the result of the experimentation is negative. Despite such complications, the paper shows that the allocations sustained under the virtual index policy are sufficiently monotone in expectation, thus permitting the construction of payments that implement the desired allocations. The application thus illustrates the broader point that, in dynamic problems, it is important not to impose ad-hoc strong monotonicity requirements on the allocation rule as a way of facilitating incentive compatibility. Many dynamic problems naturally call for allocation rules that are not strongly monotone, and yet implementable.

Two other recent papers on dynamic experimentation with private information are Kakade et al. (2013) and Fershtman and Pavan (2016). The former paper considers a dynamic allocation problem similar to the one in the bandit auction application in Pavan, Segal, and Toikka (2014). The approach, however, is slightly different. The authors confine attention to a family of problems in which payoffs are sufficiently separable (either additively, or multiplicatively) in the agents' initial types. Such separability, in turn, permits the authors to propose an implementation of the virtual index rule based on payments that are similar to those in the dynamic pivot mechanism of Bergemann and Välimäki (2010). The "trick" in the proposed mechanism is to have the agents re-report in each period their initial type. Such re-reporting "Markovizes" the environment by guaranteeing that, when the payments are the virtual VCG payments, and the allocation rule is virtually efficient both on- and off-path, each agent finds it optimal to stay in the mechanism and report truthfully at any history, as long as he/she expects all other agents to do the same. This is because, under the virtual VCG mechanism, when payoffs are sufficiently separable in the first component, each agent's continuation payoff at any history is proportional to the agent's contribution to the aggregate virtual surplus in the continuation game. Arguments similar to those in the literature on the implementation of efficient rules (see, e.g., Vickrey, 1961; Clarke, 1971; and Groves, 1973 for static settings, and Bergemann and Välimäki, 2010 and Athey and Segal, 2013 for dynamic settings) then imply that it is in each agent's own interest to induce a virtually efficient allocation in each continuation game, which can be accomplished by (a) correcting possible period-1 lies and then (b) reporting the current and future innovations truthfully.

An approach similar to the one in Kakade et al. (2013) is also used in Fershtman and Pavan (2016) in the design of dynamic matching auctions. That paper considers the problem of a platform (a matchmaker) controlling the interactions between two sides of a market. Each side is populated by agents whose value for interacting with agents from the opposite side is their own private information and evolves over time, either as the result of shocks that affect the attractiveness of individual interactions, or as the result of gradual resolution of uncertainty about fixed but unknown payoffs.[26]

[26] The paper also considers the possibility that valuations vary (stochastically) over time because of preferences for variety, by which the attractiveness of each individual match decreases with the number of past interactions.

The paper accommodates for preferences that combine elements of vertical differentiation (certain agents value interacting with all other agents more than others) with elements of horizontal differentiation (the relative attractiveness of any pair of agents from the opposite side may vary across individuals and can change over time).[27] It also accommodates for constraints on the number of possible matches that can be maintained in each period. The analysis identifies the properties of welfare- and profit-maximizing mechanisms and shows how the optimal matching dynamics can be sustained via "matching auctions." In such auctions, agents from each side first select a membership status in each period (a higher status grants preferential treatment in the auctions). They then bid for each possible partner from the opposite side. Each bilateral match is then assigned a "score" that combines the involved agents' reciprocal bids for one another, the agents' membership status, and the number of past interactions between the two agents. The matches with the highest nonnegative score are then implemented, up to capacity. The payments are designed so as to make each agent's continuation payoff proportional to the agent's contribution to the continuation weighted surplus. Depending on the designer's objective, the latter may coincide with total welfare, with the platform's profits, or, more generally, with a combination of the two. The payments are similar, in spirit, to those in the Generalized Second Price (GSP) auction used in sponsored search, but adjusted to take into account (a) the cost of leaving the agents information rents, and (b) the value of generating information that can be used in later periods. Importantly, in equilibrium, all agents find it optimal to remain in the mechanism after all histories and bid truthfully their myopic values for all partners, despite their own private value for experimentation.

The paper also compares matching dynamics under profit maximization with their counterparts under welfare maximization. When all agents derive a nonnegative value from interacting with any other agent, profit maximization involves fewer and shorter interactions than welfare maximization. In particular, when the capacity constraint is not binding, under profit maximization, each pair of agents is matched for an inefficiently short period of time. When, instead, the capacity constraint is binding, certain matches last longer under profit maximization than under welfare maximization. In each period, however, the aggregate number of interactions, is higher under welfare maximization than under profit maximization. Importantly, the paper shows that the above conclusions need not extend to markets in which certain agents dislike certain interactions. In this case, in each period, profit maximization may involve an inefficiently large volume of matches. The above conclusions have implications for government intervention in matching markets, an area that is attracting significant attention both from policy makers and market designers.

[27] Static matching design with vertically and/or horizontally differentiated preferences is examined in Gomes and Pavan (2016a, 2016b).

### 7.1 Taxation Under Learning-by-Doing

Allowing for endogenous types (that is, for decision-controlled processes governing the evolution of the agents' private information) not only opens the door to novel applications, it may also change some of the results established by confining attention to exogenous processes. This possibility is illustrated in a recent paper by Makris and Pavan (2016). They consider a dynamic model of income taxation over the life cycle, where the agents' productivity evolves endogenously over time as the result of learning-by-doing.[28]

Learning-by-doing (hereafter LBD) refers to a situation where, by working harder in the present period, a worker increases (stochastically) his future average productivity (or wages). One can think of LBD as investment in human capital that occurs "on the job," that is, through the production process (as opposed to, say, through training that occurs outside the working process).[29]

Makris and Pavan (2016) study the effects of LBD on the dynamic provision of incentives under optimal tax codes. They consider a fairly general Markovian setting with endogenous types, non-quasilinear payoffs, and (imperfect) type persistence, that embeds most models in dynamic public finance and certain models in the dynamic managerial compensation literature (e.g., the two-period model in Garrett and Pavan, 2015, discussed in the previous section[30]). In particular, the paper focuses on the effects of LBD on the dynamics of wedges (that is, distortions relative to the first-best benchmark, introduced to limit information rents). LBD affects the cost of incentives (equivalently, the expectation of future information rents) through two channels. The first one is by changing, for given future information rents, the distribution of future types. The second one is through its effect, for given distribution of future types, on future information rents via its effect on the impulse response functions. These effects are absent under an exogenous type process. They are also absent in models where the agents' productivity is endogenous, possibly because of past investment in human capital, but where the agents' private information is exogenous (e.g., Kapicka and Neira, 2014; Best and Kleven, 2013; and Stantcheva, 2016). Importantly, these effects may overturn some of the conclusions established in the literature assuming an exogenous type process. For example, the paper shows that LBD may result in wedges (equivalently,

[28] Related models of dynamic taxation with and without human capital accumulation include Krause (2009), Best and Kleven (2013), Farhi and Werning (2013), Kapicka (2013a, 2013b), Golosov, Troshkin, and Tsyvinski (2015), Kapicka and Neira (2014), Stantcheva (2016), and Heathcote et al. (2014).

[29] A vast literature in labor economics documents the effect of labor experience on wages. See, for instance, Willis (1986), Topel (1991), Jacobson et al. (1993), Altuğ and Miller (1998), and Dustmann and Meghir (2005).

[30] For a related model of dynamic managerial compensation where the principal may fire the agent and replace him with a new hire and where each new relationship is affected by shocks to managerial productivity that are privately observed by the managers, see Garrett and Pavan (2012).

marginal tax rates) that are decreasing over the life cycle, in the same environment in which the opposite dynamics obtain under an exogenous process (see, e.g., Farhi and Werning, 2013 and Garrett and Pavan, 2015).

Furthermore, LBD can contribute to higher wedges (equivalently, higher marginal taxes) at all productivity levels. Finally, LBD can contribute to a higher progressivity of the wedges, that is, to an increase in wedges that is more pronounced at the top of the current-period type distribution than at the bottom. Interestingly, in the presence of LBD, marginal taxes may be increasing in earnings in the same environments in which optimal tax codes have been shown to be regressive abstracting from LBD effects. I illustrate these possibilities below in a simplified version of the model considered in Makris and Pavan (2016).

### *7.1.1 The Environment*

The economy is the same as in the simple two-period model of Section 6, above, except for the fact that the agents' period-2 productivity is drawn from a distribution $F_2(\theta_1, y_1)$ that depends on period-1 output $y_1$. The dependence of $F_2$ on $y_1$ is what captures LBD effects. These effects in turn may originate either in past "effort/labor supply," or directly in past "output." While the notation in this section is for an economy in which $T = 2$, some of the key formulas and results will be described for arbitrary time horizons.

Denote by $F_2(\theta_2|\theta_1, y_1)$ the cdf of the kernel $F_2(\theta_1, y_1)$. For any $\theta_2 \in \Theta_2$, $F_2(\theta_2|\theta_1, y_1)$ is assumed to be nonincreasing in $(\theta_1, y_1)$. Next, let $\lambda[\chi]$ denote the endogenous ex-ante probability distribution over $\Theta$ under the rule $\chi = \langle \mathbf{y}, \mathbf{c} \rangle$ (as in the previous section, the rule $\chi$ comprises output and consumption policies that specify the evolution of the relevant allocations, as a function of the history of the agent's productivity). Finally, let $\lambda[\chi]|\theta_1, y_1$ denote the distribution over $\Theta$, starting from $(\theta_1, y_1)$, and $\lambda[\chi]|\theta_1 = \lambda[\chi]|\theta_1, y_1(\theta_1)$, where $y_1(\theta_1)$ is the output generated in period one by an agent of productivity $\theta_1$.

To make some of the formulas below easier to read, I will denote by $\mathbb{E}^{\lambda[\chi]|\theta^t}[\cdot]$ the expectation over $\Theta$, under the endogenous process $\lambda[\chi]$, starting from $\theta^t, y^t(\theta^t)$.

The principal designs the rule $\chi$ so as to maximize the ex-ante expected value of her payoff subject to the rule being incentive compatible and satisfying the constraint

$$\int q(V^A(\theta_1))dF_1(\theta_1) \geq \kappa. \tag{24}$$

As in the previous section, the function $q(\cdot)$, which is assumed to be increasing and (weakly) concave, captures the nonlinear Pareto weights the principal assigns to the agent's expected lifetime utility. This formulation captures, in reduced form, various problems considered in the literature. For example, in a taxation economy, the principal is a government maximizing tax revenues under the constraint that the agent's lifetime utility be above a minimal level (equivalently, in the dual of this problem, the government maximizes the agents' lifetime utility subject to a budget constraint). In this problem, the function $q$ captures the planner's inequality aversion (see, for instance, Farhi and Werning, 2013, and Best and Kleven, 2013).

In this section, I consider the simplest possible version of this problem where the agents are risk-neutral (using the notation from the previous section, this amounts to assuming that $v^A(c_t) = c_t$, for all $c_t$) and the principal has a Rawlsian's objective (which amounts to replacing the constraint in (24) with the simpler constraint that $V^A(\underline{\theta}_1) \geq \kappa$). I also assume that the disutility of effort takes the familiar isoelastic form

$$\psi(y_t, \theta_t) = \frac{1}{1+\phi}\left(\frac{y_t}{\theta_t}\right)^{1+\phi}, \tag{25}$$

which implies that both the elasticity of the disutility of effort and the elasticity of the marginal disutility of effort with respect to the agent's productivity are constant and equal to

$$\epsilon_\theta^\psi(y_t, \theta_t) \equiv -\frac{\theta_t \psi_\theta(y_t, \theta_t)}{\psi(y_t, \theta_t)} = 1 + \phi = \epsilon_\theta^{\psi_y}(y_t, \theta_t) \equiv -\frac{\theta_t \psi_{y\theta}(y_t, \theta_t)}{\psi_y(y_t, \theta_t)},$$

with $\phi$ parametrizing the inverse Frisch elasticity.[31] Because these restrictions are not important for the qualitative results, in the discussion below I will specialize the notation to the above functional form only when presenting some numerical results. For a more general treatment of the problem in which (a) the agents live, and work, for arbitrarily long horizons, (b) agents are risk-averse, (c) the planner has smoother aversion to inequality (captured by arbitrary $q$ functions), (d) the agents' productivity evolves according to a general type process, I refer the reader to Makris and Pavan (2016). In addition to generalizing the insights discussed below, the paper also shows how the dynamics of allocation under optimal tax codes can be derived through a recursive approach that explicitly accounts for the endogeneity of the agents' private information.

### *7.1.2 The First-Best Benchmark*

In order to appreciate how LBD affects the level, progressivity, and dynamics of the wedges, consider first the allocations that solve the principal's problem, in the absence of information frictions, when the process governing the evolution of the agents' productivity is endogenous. Let

$$LD^{FB;\chi}(\theta_1) \equiv \delta \frac{\partial}{\partial y_1} \int \{y_2(\theta) - \psi(y_2(\theta), \theta_2)\} \, dF(\theta_2|\theta_1, y_1(\theta_1)) \tag{26}$$

denote the effect of a marginal change in period-1 output on the expected sum of the principal's and the agent's continuation payoffs, under the policy $\chi$, evaluated at history $(\theta_1, y_1(\theta_1))$.

[31] As in the previous section, the notation $\psi_\theta(y_t, \theta_t)$ and $\psi_y(y_t, \theta_t)$ stands for the partial derivative of the $\psi$ function with respect to $\theta_t$ and $y_t$, respectively, whereas $\psi_{y\theta}(y_t, \theta_t)$ stands for the cross derivative.

**Proposition 2 (Makris and Pavan, 2016)** *The first-best rule* $\chi^* = \langle \mathbf{y}^*, \mathbf{c}^* \rangle$ *satisfies the following optimality conditions (at all interior points with* $\lambda[\chi^*]$*-probability one)*[32]

$$1 + LD^{FB;\chi^*}(\theta_1) = \psi_y(y_1^*(\theta), \theta_1) \text{ and } 1 = \psi_y(y_2^*(\theta), \theta_2). \tag{27}$$

The principal thus equalizes the marginal benefit of asking the agent for higher output (taking into account, in period one, the effect from LBD) with its marginal cost. When the agent is risk-neutral, the latter is simply the cost of increasing the agent's compensation by an amount equal to the agent's marginal disutility of effort. As to the marginal benefit, observe that the sum of the principal's and of the agent's continuation payoffs is increasing in $\theta_2$ (recall that $\psi_y(y_2, \theta_2)$ is decreasing in $\theta_2$). Because LBD shifts the distribution in a first-order-stochastic-dominance way, $LD^{FB;\chi^*}(\theta_1) \geq 0$. Hence, LBD naturally induces the principal to ask for a higher output in period 1 compared to the level she would ask in the absence of LBD.

As shown in Makris and Pavan (2016), under risk aversion, the above optimality conditions must be paired with other optimality conditions that require the equalization of the marginal utility of consumption between any two consecutive histories. Furthermore, with more general objectives (that is, with less extreme preferences for redistribution than in the Rawlsian case assumed here), optimality also requires equalizing across period-1 types the "weights" the principal assigns to the agent's period-1 marginal utility of consumption.

### *7.1.3 Dynamics of Wedges (or Marginal Distortions) under Second-Best Policies*

In contrast to the previous section, here I follow the First-Order Approach (FOA) by considering a *relaxed program* in which the incentive compatibility constraints are replaced by the envelope conditions requiring that, for all $t$, all $\theta^{t-1}$, the agents' equilibrium continuation payoff under $\chi$ be Lipschitz continuous in $\theta_t$, with derivative equal, for almost all $\theta_t \in \Theta_t$, to

$$\frac{\partial V_t^A(\theta^t)}{\partial \theta_t} = -\mathbb{E}^{\lambda[\chi]|\theta^t}\left[\sum_{\tau=t}^{T} \delta^{\tau-t} I_{(t),\tau}(\theta^\tau, y^{\tau-1}(\theta^{\tau-1}))\psi_\theta(y_\tau(\theta^\tau), \theta_\tau)\right]. \tag{28}$$

The discussion of the various channels through which LBD affects the dynamics of distortions below should therefore be understood as being conditional on the First-Order Approach being valid. Note that the assumption that each $F_t(\theta_t|\theta_{t-1}, y_{t-1})$ is nonincreasing in $\theta_{t-1}$ implies that the impulse responses are nonnegative, and hence that each agent's expected lifetime utility is nondecreasing in $\theta_t$. In particular, that $V_1^A(\theta_1)$ is nondecreasing in $\theta_1$ implies that

[32] Here I focus on the optimality conditions for the output schedule. The optimality conditions for consumption are discussed in Makris and Pavan (2016).

the taxation problem with a risk-neutral agent and a planner with a Rawlsian objective under examination here is formally equivalent to a managerial-compensation problem with interim individual-rationality constraints (where ensuring the participation constraint of the period-1 lowest types implies that all other types' participation constraints are also satisfied, as discussed above).

As in the previous section, to limit the agents' information rents, the principal distorts production downwards. In the presence of LBD, the distortions in labor supply are captured by labor wedges, defined as follows:

**Definition 8 (wedges with LBD)** *The period-t "wedge" or "marginal distortion" (equivalently, the "effective marginal income tax rate" in the taxation problem) under the rule $\chi$ is given by* [33]

$$W_1(\theta) \equiv 1 - \frac{\psi_y(y_1(\theta), \theta_1)}{1 + LD^{FB;\chi}(\theta_1)} \text{ and } W_2(\theta) = 1 - \psi_y(y_2(\theta), \theta_2).$$

The definition parallels the one in the previous section, adjusted for the presence of LBD. Recall that efficiency requires that the marginal disutility of extra period-$t$ output be equalized to the social marginal benefit, where, in period one, the latter takes into account both the extra output collected by the principal, 1, and the effect of higher period-1 output on the principal's and the agent's joint future surplus, as captured by the term $LD^{FB;\chi}(\theta_1)$ introduced above. The wedge is the discrepancy between the ratio of marginal cost and marginal benefit of higher period-$t$ output under the first-best allocation, 1, and the corresponding ratio at the proposed allocation. Importantly, in the case of an endogenous process, such discrepancy is computed holding fixed the rule $\chi$ that determines future allocations, so as to highlight the part of the inefficiency that pertains to the period under consideration. As is standard in the public finance literature, hereafter I will focus on the dynamics of the following monotone transformation of the wedges

$$\hat{W}_t(\theta) \equiv \frac{W_t(\theta)}{1 - W_t(\theta)}$$

and refer to $\hat{W}_t(\theta)$ as to the *relative period-t wedge*. In a taxation problem, the wedge is directly related to the period-$t$ marginal tax rate; that is, the sensitivity of current taxes to current income, holding fixed past incomes and all future tax schedules (with the latter allowed to depend on the entire history of reported earnings). One can easily verify that the relative period-$t$ wedge $\hat{W}_t(\theta)$ satisfies

$$1 + LD_t^{FB;\chi}(\theta) = \left[1 + \hat{W}_t(\theta)\right] \psi_y(y_t(\theta), \theta_t). \tag{29}$$

[33] Obviously, $y_1$ and $W_1$ naturally depend on $\theta = (\theta_1, \theta_2)$ only through $\theta_1$. The reasons why these functions, in the formulas below, are allowed to depend on the entire type history $\theta$ is just to economize on notation by introducing homogeneity in the arguments across different periods.

One can thus also think of the relative wedge as the rate by which the principal inflates the agent's marginal disutility of labor to account for the effects that higher output has on informational rents.

Finally, let

$$h_2(\theta, y) \equiv \frac{1 - F_1(\theta_1)}{f_1(\theta_1)} I_2(\theta, y_1)\epsilon_\theta^\psi(y_2, \theta_2)\frac{\psi(y_2, \theta_2)}{\theta_2}.$$

This function is the analog in the taxation environment under consideration here of the "handicap" function $I_2(\theta)\psi_y(y_2, \theta_2)\left[1 - F_1(\theta_1)\right]/f_1(\theta_1)$ in the managerial-compensation model in the previous section. It measures the welfare losses of asking the agent to produce higher output at history $\theta = (\theta_1, \theta_2)$.

The second-best allocation rule $\chi$ satisfies the following optimality conditions (at all interior points with $\lambda[\chi]$-probability one):

$$1 = \psi_y(y_2(\theta), \theta_2)\left[1 + \frac{1 - F_1(\theta_1)}{f_1(\theta_1)} I_2(\theta, y_1(\theta_1))\frac{\epsilon_\theta^{\psi_y}(y_2(\theta), \theta_2)}{\theta_2}\right] \tag{30}$$

$$1 + LD^{FB;\chi}(\theta_1) = \psi_y(y_1(\theta), \theta_1) \times \left[1 + \frac{1 - F_1(\theta_1)}{f_1(\theta_1)}\frac{\epsilon_\theta^{\psi_y}(y_1(\theta), \theta_1)}{\theta_1} + \delta\frac{\partial}{\partial y_1}\mathbb{E}^{\lambda[\chi]|\theta_1, y_1(\theta)}[h_2(\theta, y(\theta))]\right] \tag{31}$$

The first optimality condition is essentially the same as in Mirrlees (1971), but adjusted for the impulse responses. At each period-2 history $\theta = (\theta_1, \theta_2)$, the optimal choice of $y_2(\theta)$ is obtained by equalizing the marginal benefit of asking the agent for higher output, the left-hand side in (30), with the marginal cost. The latter combines the cost of reimbursing the agent for the higher disutility of effort, $\psi_y(y_2(\theta), \theta_2)$, with the marginal cost of increasing the rents left to all period-1 types whose productivity is above $\theta_1$. The latter cost is given by

$$\frac{\partial}{\partial y_2}h_2(\theta, y(\theta)) = \frac{1 - F_1(\theta_1)}{f_1(\theta_1)} I_2(\theta, y_1(\theta_1))\frac{\epsilon_\theta^{\psi_y}(y_2(\theta), \theta_2)}{\theta_2}\psi_y(y_2(\theta), \theta_2).$$

The optimality condition in (31), instead, combines the above static effects with the dynamic effects that a higher $y_1$ has on the principal's payoff, net of the agents' informational rents. As discussed above, holding fixed the period-2 policies, as specified in $\chi = \langle \mathbf{y}, \mathbf{c} \rangle$, the term $LD^{FB;\chi}(\theta_1)$ captures the effect on total surplus of shifting the agent's period-2 productivity towards higher levels. The interesting novel effects here are those captured by the term

$$\frac{\partial}{\partial y_1}\mathbb{E}^{\lambda[\chi]|\theta_1, y_1(\theta)}[h_2(\theta, y(\theta))] = \frac{1 - F_1(\theta_1)}{f_1(\theta_1)}\frac{\partial}{\partial y_1}\mathbb{E}^{\lambda[\chi]|\theta_1, y_1(\theta)} \times \left[I_2(\theta, y_1(\theta_1))\frac{\epsilon_\theta^{\psi_y}(y_2(\theta), \theta_2)}{\theta_2}\psi_y(y_2(\theta), \theta_2)\right].$$

Recall that the term $h_2(\theta, y(\theta))$ captures the welfare losses coming from the information rents the principal must leave to the agents to induce them to report truthfully in period 2. Asking the agents to work harder in period one has two effects on such expected losses. The first one comes from the change in the distribution of $\theta_2$, holding the handicap function $h_2(\theta, y(\theta))$ constant. The second comes from the direct effect that a higher period-1 output has on the period-2 handicap $h_2(\theta, y(\theta))$ via the endogenous impulse response of $\theta_2$ to $\theta_1$, holding the period-2 type distribution constant.

In most cases of interest, the period-2 handicap (equivalently, the period-2 information rent) $h_2(\theta, y(\theta))$ is increasing in $\theta_2$. In this case, the first effect contributes to higher expected welfare losses. The sign of the second effect in turn depends on whether the impulse response of $\theta_2$ to $\theta_1$ is increasing or decreasing in $y_1$. When period-1 type and period-1 output are complements in the determination of the period-2 type, impulse responses are increasing in $y_1$, in which case the second effect also contributes to a positive effect of LBD on expected future welfare losses. The opposite is true when period-1 type and period-1 output are substitutes in the determination of the period-2 type. In this case, this second effect operates against the first one in alleviating the positive effects of LBD on expected future rents.

The above results can be used to interpret the dynamics of wedges under optimal contracts. To this purpose, let

$$\hat{W}_t^{RN}(\theta^t) \equiv \frac{1 - F_1(\theta_1)}{f_1(\theta_1)} I_t(\theta^t, y^{t-1}(\theta^{t-1})) \frac{\epsilon_\theta^{\psi_y}(y_t(\theta^t), \theta_t)}{\theta_t}$$

and

$$\Omega(\theta_1) \equiv \frac{\delta \frac{\partial}{\partial y_1} \mathbb{E}^{\lambda[\chi]|\theta_1, y_1(\theta)}[h_2(\theta, y(\theta))]}{\psi_y(y_1(\theta_1), \theta_1)}.$$

From (29), (30) and (31), one can see that

$$\hat{W}_1(\theta_1) = \hat{W}_1^{RN}(\theta_1) + \Omega_1(\theta_1) \text{ and } \hat{W}_2(\theta) = \hat{W}_2^{RN}(\theta).$$

The functions $\hat{W}_t^{RN}(\theta^t)$ represent the relative wedges that the principal would select at each type history $\theta^t$ if the process governing the evolution of the agents' productivity determined by the rule $\chi$ were exogenous (when the agents are risk-neutral and the principal's objective is Rawlsian, as assumed here). Because there are no LBD effects in the last period (here $t = 2$), naturally, $\hat{W}_2(\theta) = \hat{W}_2^{RN}(\theta)$. The function $\Omega_1(\theta_1)$, instead, captures the effects of LBD on the period-1 wedge. As explained above, it summarizes all the effects of changing the period-1 output on the expected informational rents the principal must leave to the agents in period 2 to induce them to report $\theta_2$ truthfully.

Makris and Pavan (2016) show that the above decomposition holds more generally; that is, it extends to environments in which the agents are risk-averse, the principal's inequality aversion is less than Rawlsian, and there are arbitrarily many periods. In particular, continue to denote by $\hat{W}_t^{RN}(\theta^t)$ and $\Omega_t(\theta^t)$, respectively, the relative wedge in the absence of LBD and the effect of LBD on the period-t wedge in the case the agents are risk-neutral and the principal's objective is Rawlsian. Then let $RA_t(\theta^t)$ be a correction that applies when the agents are risk-averse and/or the principal's inequality aversion is less than Rawlsian.

**Theorem 7 (Makris and Pavan, 2016)** *At any period $t \geq 1$, with $\lambda[\chi]$-probability one, the relative wedges are given by*

$$\hat{W}_t(\theta^t) = RA_t(\theta^t)\left[\hat{W}_t^{RN}(\theta^t) + \Omega_t(\theta^t)\right].$$

Theorem 7 is established using a recursive solution to the principal's problem that controls for the endogeneity of the type process. The value of the theorem is in highlighting the different forces that contribute to the determination of the dynamics of the wedges. The theorem also unifies the various special cases examined in the literature.

To see how the endogeneity of the type process shapes the dynamics of the wedges in the simplest possible way, consider the following specification often assumed in the taxation and in the labor economics literature. The period-2 productivity is given by

$$\theta_2 = z_2(\theta_1, y_1, \varepsilon_2) = \theta_1 y_1^{\zeta} \varepsilon_2, \tag{32}$$

where $\zeta \geq 0$ parametrizes the intensity of LBD (the case of no LBD corresponds to $\zeta = 0$) and where $\varepsilon_2$ is a shock drawn from a distribution $G$. The specification in (32) implies that the impulse responses of $\theta_2$ to $\theta_1$ are invariant in $y_1$ and are given by $I_1(\theta) = 1$ and $I_2(\theta) = \theta_2/\theta_1$, all $\theta = (\theta_1, \theta_2)$. This specification has the advantage of isolating the effects of LBD coming from the shift in the distribution of period-2 productivity (the other channel by which LBD affects the expectation of the period-2 rents via its direct effect on the impulse response functions is mute under the above specification). A second advantage of this specification is that it implies that the period-2 wedges are invariant in $\zeta$, which facilitates the comparison to the case without LBD. In particular, one can use the results above to verify that, when the disutility of effort takes the isoelastic form in (25) and LBD takes the multiplicative form in (32),

$$\hat{W}_1^{RN}(\theta_1) = \hat{W}_2^{RN}(\theta) = (1+\phi)\,\frac{1-F_1(\theta_1)}{f_1(\theta_1)\theta_1}$$

and

$$\Omega_1(\theta_1) = \delta\left(\frac{\theta_1}{y_1(\theta_1)}\right)^{\phi}\frac{1-F_1(\theta_1)}{f_1(\theta_1)}\frac{\partial}{\partial y_1}\mathbb{E}^{\lambda[\chi]|\theta_1, y_1(\theta_1)}\left[\left(\frac{y_2(\theta)}{\theta_2}\right)^{1+\phi}\right]. \tag{33}$$

Hence, when $[1 - F_1(\theta_1)]/\theta_1 f_1(\theta_1)$ is nonincreasing in $\theta_1$, as typically assumed in the taxation literature, $\hat{W}_1^{RN}(\theta_1)$ and $\hat{W}_2^{RN}(\theta)$ are nonincreasing in $\theta_1$ and independent of $\theta_2$. That is, in the absence of LBD, marginal taxes are constant over time and decreasing in earnings. Instead, with LBD, wedges are given by

$$\hat{W}_1(\theta^t) = \hat{W}_1^{RN}(\theta_1) + \Omega_1(\theta_1) \text{ and } \hat{W}_2(\theta) = \hat{W}_2^{RN}(\theta).$$

As shown in Makris and Pavan (2016), in this economy, as well as in most cases considered in the literature, the period-2 handicap is increasing in $\theta_2$ (equivalently, $y_2(\theta)/\theta_2$ is increasing in $\theta_2$). As a result, $\Omega_1(\theta_1) > 0$. The following is then true:

**Proposition 3 (Makris and Pavan, 2016)** *Consider the two-period economy described above. (i) For all $\theta_1$, the period-1 wedge in the presence of LBD is strictly higher than the corresponding period-1 wedge in the absence of LBD; (ii) Under LBD, in each state of the world, period-1 wedges are strictly higher than period-2 wedges; (iii) There exits a function $\Gamma : \Theta_1 \to \mathbb{R}$ such that the progressivity of the period-1 wedge at productivity level $\theta_1$ is higher under LBD than in its absence if and only if* $\Gamma(\theta_1) \geq 0$.

As an illustration, suppose $\theta_1$ is drawn from a Pareto distribution (in which case, $\theta_1 f_1(\theta_1)/[1 - F_1(\theta_1)] = \lambda$ for all $\theta_1$), as assumed in the earlier taxation literature. In the absence of LBD, the wedges are constant over time and across types and are given by $\hat{W}_1(\theta_1) = \hat{W}_2(\theta) = (1 + \phi)/\lambda$, for all $\theta = (\theta_1, \theta_2)$. Instead, in the presence of LBD, $\hat{W}_1(\theta_1) > \hat{W}_2(\theta) = (1 + \phi)/\lambda$ with $\hat{W}_1(\theta_1)$ strictly increasing in $\theta_1$. Furthermore, one can verify that, whenever the period-1 wedge in the absence of LBD, $\hat{W}_1^{RN}(\theta_1)$, is nonincreasing in $\theta_1$, as is the case here, the solution to the relaxed program also solves the full program, implying that the properties above are truly features of optimal tax codes. This situation is illustrated in Figure 4. The figure depicts the period-1 wedge $\hat{W}_1(\theta_1)$ for four different levels of LBD, $\zeta = 0, 0.2, 0.4, 0.6$. As the figure illustrates, stronger LBD effects (here captured by higher levels of the parameter $\zeta$) call for higher and more progressive wedges.[34]

The intuition is the one anticipated above. LBD affects the cost of future incentives (equivalently, expected future information rents) through two channels. The first one is by changing the distribution of future types, for given future handicaps (recall that the latter measure future welfare losses due to asymmetric information). The second channel is through its direct effect on future handicaps for given distribution of future types. As mentioned above,

[34] The parameter $eta = 0$ in the Figure's caption indicates that the results are for the case of risk-neutral agents. Also note that the figure assumes $\phi = 2$, i.e., a Frisch elasticity of 0.5, as in Farhi and Werning (2013), Kapicka (2013a), and Stantcheva (2016). Finally, the parameter $rho = 1$ in the figure's caption indicates that the skill-persistence parameter has been set equal to one, as in the rest of this section (i.e., $\theta_2 = \theta_1^{\rho} y_1^{z} \varepsilon$, with $\rho = 1$).

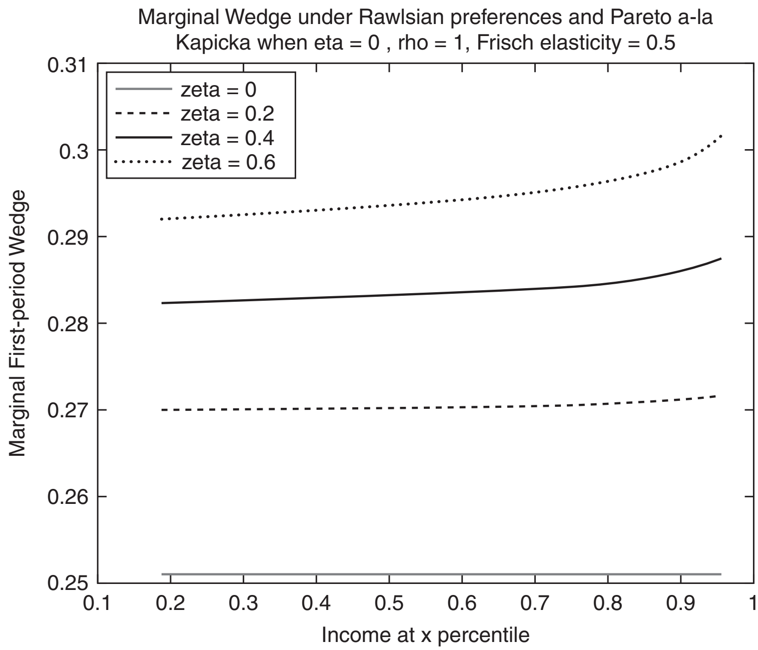


Figure 4 The Risk-Neutral, Rawlsian, Pareto case.

this second channel is absent under the specification in (32). More generally, this second channel adds to the first one in amplifying the effects of LBD on expected future rents when impulse responses of future types to current ones are increasing in output.

To reduce the effects of LBD on expected future rents, the principal induces the agents to work less in the early periods of their career, which explains why LBD contributes to higher wedges. Because the effects of LBD on expected future rents are most pronounced in the early periods, LBD also contributes to wedge dynamics whereby wedges decline over time.

Finally, to see why LBD may contribute to the progressivity of the wedges, note that the benefit of distorting downwards labor supply in the present period so as to economize on future information rents is stronger for higher types, given that these are the types that expect, on average, larger rents in future periods. This is indeed the case, for all $\theta_1$, when $\theta_1$ is drawn from a Pareto distribution. More generally, Makris and Pavan (2016) show that this is true, in the simple economy under consideration here, at any open interval of period-1 types at which $\theta_1 f_1(\theta_1)/[1 - F_1(\theta_1)]$ is almost constant, as in the upper tail of the Pareto Lognormal distribution typically assumed to calibrate the model to US earnings data.

The next four figures illustrate that the properties discussed above hold more generally. In particular, Figure 5 illustrates the effects of LBD on the period-1 wedge $\hat{W}_1$ when $\theta_1$ and the second-period shock $\varepsilon_2$ are drawn

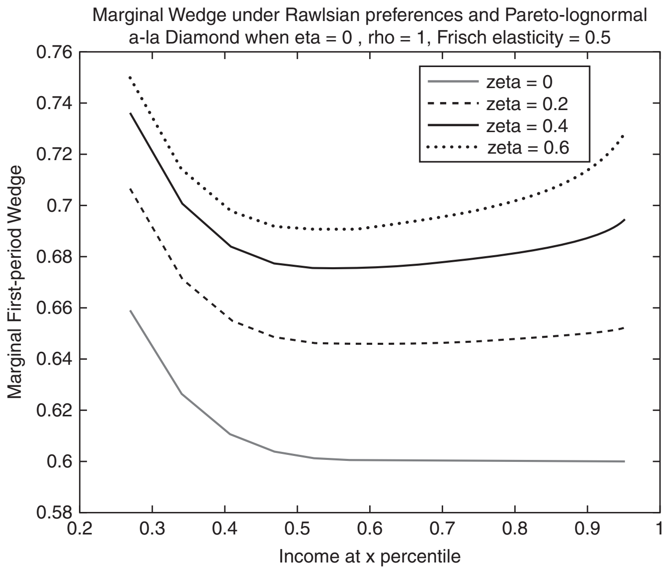


Figure 5 The Risk-Neutral-Rawlsian-Paretolognormal case.

from a Pareto Lognormal distribution with the same parameters as in Diamond (1998). As the figure shows, stronger LBD effects (here parametrized by a higher level of $\zeta$) are responsible for higher period-1 wedges and for more progressivity at all income percentiles, and in particular at high percentiles (the lowest blue curve corresponds to the economy without LBD of Diamond, 1998).

Similar results obtain when the agents are risk-averse and/or the planner's aversion to inequality is less extreme than in the Rawlsian case. Figure 6 depicts the wedges in the case the agents' utility over consumption is CRRA for four different levels of the coefficient of relative risk aversion, namely for $\eta = 0$, 0.2, 0.5, and 0.8 (recall that $\eta = 0$ corresponds to the case of risk-neutral agents). The distribution of $\theta_1$ and $\varepsilon_2$ is Pareto-lognormal and the principal's objective is Rawlsian, as in Figure 5 (higher curves correspond to higher degrees of risk aversion, whereas higher levels of the zeta parameter correspond to stronger LBD effects).

As the figure illustrates, risk aversion contributes to an amplification of the period-1 wedge and to more progressivity. The reason is that, when the agent is risk-averse, the cost of compensating him for his marginal disutility of effort is higher as it now takes $\psi_y(\theta_t, y_t)/v'(c_t)$ units of consumption. Hence, risk aversion, by itself, contributes to higher wedges. What the figure illustrates is that the increase in the cost of compensation also amplifies the effects of LBD on the level and progressivity of the wedges. This is because, by increasing the cost of future information rents, risk aversion also increases the benefits of shifting the distribution of future types towards levels that command lower informational rents, and, as in the risk-neutral case, this effect is most

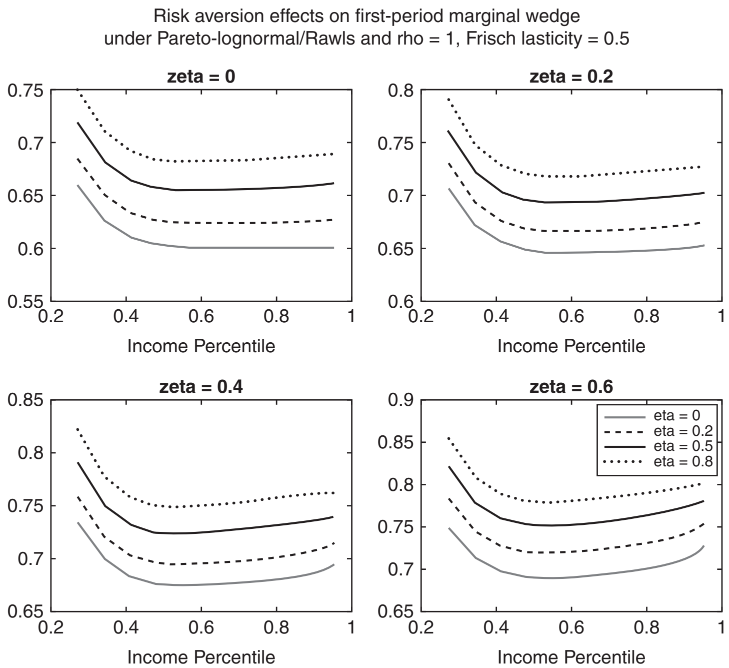


Figure 6 The CRRA Rawlsian Pareto Lognormal case.

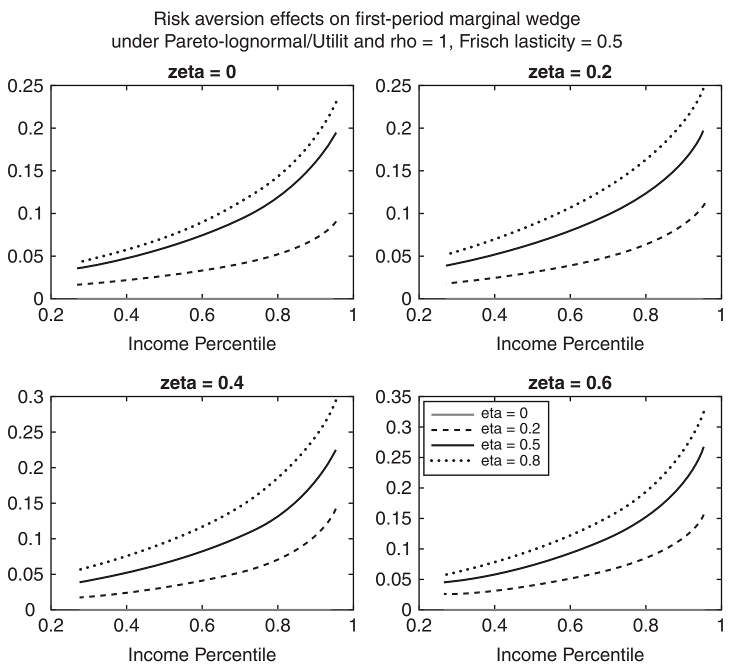


Figure 7 The CRRA Utilitarian Pareto Lognormal case.

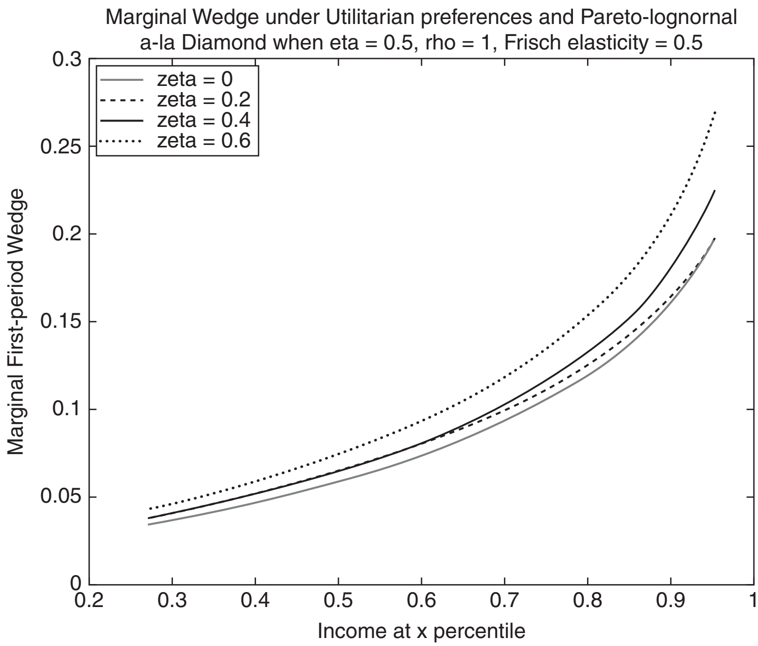


Figure 8 Effects of LBD in CRRA Utilitarian Pareto Lognormal case.

pronounced at the top of the period-1 type distribution where the expectation of future rents is the highest.

Similar results obtain for less extreme forms of inequality aversion on the planner's side. Figure 7 reports the results for the case of a planner with utilitarian objective function – the parameters are as in Figure 6. Once again, risk aversion contributes to an amplification of the effects discussed above, at all levels of LBD. Furthermore, fixing the degree of relative risk aversion to $\eta = 0.5$, stronger LBD effects contribute to higher wedges and to more progressivity, as highlighted in Figure 8 (higher curves correspond to stronger LBD effects).

## 8 CONCLUSIONS

Dynamic mechanism design has proved useful to conduct positive and normative analysis in environments in which information arrives over time and a sequence of decisions is to be made. The approach pioneered by Myerson (1981) for static design problems can be extended to dynamic environments. An important tool of this approach is an envelope representation of the agents' equilibrium payoffs summarizing necessary local incentive compatibility conditions. The key difference with respect to static problems is the presence in the envelope formula of "impulse response functions" capturing the marginal effects of current types on future ones. In quasilinear settings, the envelope formula pins down the transfers, for given non-monetary allocation rule. In Markov environments, this formula, jointly with appropriate integral monotonicity conditions on the non-monetary allocation rule, provides a complete characterization of incentive compatibility (Pavan, Segal, and Toikka, 2014).

The approach typically followed in applications to solve for optimal mechanisms consists in disregarding the integral monotonicity conditions and solving a relaxed program in which only the necessary conditions captured by the envelope formulas of the equilibrium payoffs are imposed. This approach, often referred to as the First-Order-Approach (FOA), offers the convenience of bypassing the difficulty of dealing with ironing and bunching, which are notoriously cumbersome even in static settings.

An important question for this literature is the extent to which the predictions identified by restricting attention to settings in which the FOA is valid extend to broader settings. An approach recently introduced in Garrett and Pavan (2015) consists in identifying perturbations of the non-monetary policies that alter payoffs while preserving integral monotonicity. When applied to policies that are incentive compatible, such perturbations preserve incentive compatibility. The perturbations can then be used to identify certain properties of the dynamics of allocations under optimal contracts, without solving for the optimal contracts. For example, in the dynamic managerial compensation model of Garrett and Pavan (2015), such a variational approach permits one to study how risk aversion and type persistence jointly interact in shaping the dynamics of average distortions. In ongoing work, Garrett, Pavan, and Toikka (2016) apply a similar variational approach to study the long-run average dynamics of optimal screening contracts.

Most of the dynamic mechanism design literature assumes the agents' private information is exogenous. In many problems of interest, however, types are endogenous. For example, valuations are endogenous in experimentation settings, and productivity is endogenous in the presence of learning-by-doing. Endogenous processes bring novel effects that qualify (and, in certain cases, overturn) the conclusions obtained assuming exogenous types. For example, in a taxation environment, learning-by-doing may change the predictions the theory delivers for the level, the dynamics, and the progressivity of the marginal tax rates, under optimal tax codes (Makris and Pavan, 2016).

A lot of work remains to be done. One direction is to extend the analysis to settings with partial commitment and/or realistic political economy constraints. The literature on limited commitment has made important progress in recent years (see, for example, Skreta, 2015; Liu et al., 2015; Maestri, 2016; Strulovici, 2016; and Gerardi and Maestri, 2016). However, this literature assumes information is static, thus abstracting from the questions at the heart of the dynamic mechanism design literature. I expect interesting new developments to come out from combining the two literatures.

Most of the dynamic mechanism design literature assumes time-varying information, but a constant population. A conspicuous literature in Operation

Research assumes constant information but a dynamic population. This latter OR literature has been extended in recent years to accommodate for private information (see, for example, the recent manual by Gershkov and Moldovanu, 2014). Combining time-varying information with population dynamics is likely to bring new insights on the structure of optimal contracts and on the dynamics of distortions due to asymmetric information (see Garrett, 2016a, 2016b).

The last few years have also witnessed a renewed interest in information design (see, e.g., Bergemann and Välimäki, 2006; Calzolari and Pavan, 2006a, 2006b; Eső and Szentes, 2007; Li and Shi, 2015; and Bergemann and Wambach, 2015 for models of endogenous disclosure in mechanism design, and Kamenica and Gentzkow, 2011; and Ely et al., 2016, among others, for persuasion in other strategic environments; see also Hörner and Skrzypacz, 2016, for an excellent overview of this literature). I expect interesting results to emerge from dynamic mechanism design models in which part of the endogeneity of the information comes from persuasion.

Lastly, most of the recent developments in the dynamic mechanism design literature have been primarily theoretical. Rich data sets have become available in the last few years about various dynamic contracting environments (insurance, health, consumer recognition, etc.). Most of the empirical contracting literature has confined attention to static problems and only in recent times has started examining dynamic settings (see, for example, Handel, Hendel, and Whinston, 2015; and Einav et al., 2010, 2013). The empirical literature, however, assumes stationary private information. I expect interesting developments will come from this empirical literature as soon as it starts accommodating for time-varying private information.

## References

Abdulkadiroğlu, A., and Loertscher, S., (2007). "Dynamic House Allocations," working paper, Duke University and University of Melbourne.

Acemoglu, D., Golosov, M., and Tsyvinski, A., (2011). "Political Economy of Ramsey Taxation." *J. Public Econ.* 95, 467–75.

Aghion P., Fudenberg D., -Holden R., Kunimoto T., and Tercieux. O., (2012). "Subgame Perfect Implementation Under Information Perturbations," working paper, Harvard University.

Akan, M., Ata, B., and Dana, J., (2015). "Revenue Management by Sequential Screening." *J. Econ. Theory* 159 (Part B), 728–74.

Albanesi, S., and Sleet, C., (2006). "Dynamic Optimal Taxation with Private Information." *Rev. Econ. Stud.* 73, 1–30.

Altuğ, S., and Miller, R. (1998). "The Effect of Work Experience on Female Wages and Labour Supply." *The Review of Economic Studies*, 65(1), 45–85.

Alvarez, F., and Jermann, U., (2000). "Efficiency, Equilibrium, and Asset Pricing with Risk of Default." *Econometrica* 68, 775–97.

Armstrong, M., and Zhou, J., forthcoming. "Search Deterrence." *Rev. Econ. Stud.*

Aron-Dine, A., Cullen, M., Einav, L., and Finkelstein, A., forthcoming. "Moral Hazard in Health Insurance: Do Dynamic Incentives Matter?" *Rev. Econ. Stat.*

Athey, S., and Segal, I., (2013). “An Efficient Dynamic Mechanism.” *Econometrica* 81, 2463–85.

Atkenson, A., and Lucas, R., (1992). “On the Efficient Distribution with Private Information.” *Rev. Econ. Stud.* 59, 427–54.

Ausubel, L. M., and Deneckere, R. J., (1989). “A Direct Mechanism Characterization of Sequential Bargaining with One-Sided Incomplete Information.” *Journal of Economic Theory* 48(1), 18–46.

Ausubel, L. M., and Deneckere, R. J., (1993). “Efficient Sequential Bargaining.” *The Review of Economic Studies* 60(2), 435–61.

Baron, D., and Besanko, D., (1984). “Regulation and Information in a Continuing Relationship.” *Inf. Econ. Policy* 1, 267–302.

Baron, D., and Myerson, R., (1982). “Regulating a Monopolist with Unknown Costs.” *Econometrica* 50, 911–30.

Battaglini, M., (2005). “Long-Term Contracting with Markovian Consumers.” *Am. Econ. Rev.* 95, 637–58.

Battaglini, M., (2007). “Optimality and Renegotiation in Dynamic Contracting.” *Games Econ. Behav.* 60, 213–46.

Battaglini, M., and Lamba, R., (2015). “Optimal Dynamic Contracting: The First-Order Approach and Beyond.” Discussion paper, Cornell University and Penn State University.

Bergemann, D., and Hörner, J., (2014). “Should Auctions be Transparent?” Discussion paper 1764R. Cowles Foundation for Research in Economics, Yale University.

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

Bergemann, D. and Morris, S., (2016). “Information Design, Bayesian Persuasion, and Bayes Correlated Equilibrium,” *Am. Econ. Rev.*, 106, no. 5.

Bergemann, D., and Pavan, A., (2015). “Introduction to JET Symposium on Dynamic Contracts and Mechanism Design.” *J. Econ.Theory* 159, 679–701.

Bergemann, D., and Pesendorfer, M., (2007). “Information Structures in Optimal Auctions.” *J. Econ. Theory* 137, 580–609.

Bergemann, D., and Said, M., (2011). “Dynamic Auctions: A Survey.” In: Cochran, J. (Ed.), *Wiley Encyclopedia of Operations Research and Management Science*. Wiley, New York, pp. 1511–22, no. 1757.

Bergemann, D., and Strack, P., (2015a). “Dynamic Revenue Maximization: A Continuous Time Approach.” *J. Econ. Theory 159* (Part B), 819–53.

Bergemann, D., and Strack, P., (2015b). “Stationary Dynamic Contracts.” Discussion paper. Yale University and UC Berkeley.

Bergemann, D., and Välimäki, J., (2002). “Information Acquisition and Efficient Mechanism Design.” *Econometrica* 70, 1007–33.

Bergemann, D., and Välimäki, J., (2006). “Information in Mechanism Design.” In: Blundell, R., Newey, W., Persson, T. (Eds.), *Advances in Economics and Econometrics*. Cambridge: Cambridge University Press, pp. 186–221.

Bergemann, D., and Välimäki, J., (2010). “The Dynamic Pivot Mechanism.” *Econometrica* 78, 771–90.

Bergemann, D., and Wambach, A., (2015). “Sequential Information Disclosure in Auctions.” *J. Econ. Theory* 159 (Part B), 1074–95.

Berger, A., Müller, R., and Naeemi, S., (2010). “Path-Monotonicity and Truthful Implementation.” Discussion paper. Maastricht University.

Besanko, D., (1985). "Multi-Period Contracts Between Principal and Agent with Adverse Selection." *Econ. Lett.* 17, 33–7.

Best, M.C., and Kleven, H.J., (2013). "Optimal Income Taxation with Career Effects of Work Effort." mimeo, LSE.

Bester, H., and Strausz, R., (2000). "Imperfect Commitment and the Revelation Principle: The Multi-Agent Case." *Econ. Lett.* 69, 165–71.

Bester, H., and Strausz, R., (2001). "Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case." *Econometrica* 69, 1077–98.

Billingsley, P., (1995). *Probability and Measure*. Wiley Series in Probability and Mathematical Statistics.

Board, S., (2007). "Selling Options." *J. Econ. Theory* 136, 324–40.

Board, S., (2011). "Relational Contracts and the Value of Loyalty." *Am. Econ. Rev.* 101, 3349–67.

Board, S., and Skrzypacz, A., (2016). "Revenue Management with Forward-Looking Buyers." *J. Polit. Econ.* 124(6) 1515–62.

Boleslavsky, R., and Said, M., (2013). "Progressive Screening: Long-Term Contracting with a Privately Known Stochastic Process." *Rev. Econ. Stud.* 80, 1–34.

Börgers, T., (2015). *An Introduction to the Theory of Mechanism Design*. Oxford: Oxford University Press.

Bulow, J., (1982). "Durable Good Monopolists." *J. Polit. Econ.* 90, 314–32.

Burguet, R., and Sákovics, J., (1996). "Reserve Prices Without Commitment." *Games Econ. Behav.* 15, 149–64.

Caillaud, B., and Mezzetti, C., (2004). "Equilibrium Reserve Prices in Sequential Ascending Auctions." *J. Econ. Theory* 117, 78–95.

Calzolari, G., and Pavan, A., (2006a). "Monopoly with Resale." *Rand J. Econ.* 37, 362–75.

Calzolari, G., and Pavan, A., (2006b). "On the Optimality of Privacy in Sequential Contracting." *J. Econ. Theory* 130, 168–204.

Carbajal, J., and Ely, J., (2013). "Mechanism Design Without Revenue Equivalence." *J. Econ. Theory* 148, 104–33.

Carroll, G., (2015). "Robustness and Linear Contracts," *American Economic Review* 105 (2), 536–63.

Carroll, G., and Meng, D., (2016). "Robust Contracting with Additive Noise." Discussion Paper, Northwestern University.

Chassang, S., (2013). "Calibrated Incentive Contracts," *Econometrica*, 81(5), 1935–71.

Che, Y., and Hörner, J., (2015). "Optimal Design for Social Learning." Discussion paper. Columbia University and Yale University.

Chiappori, P.-A., and Salanié, B., (2003). "Testing Contract Theory: A Survey of Some Recent Work." In: Dewatripont, M., Hansen, L., Turnovsky, S. (eds.), *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*, Vol. 1. Cambridge University Press, pp. 115–49.

Chiappori, P., Salanié, B., and Valentin, J., (1999). "Early Starters Versus Late Beginners." *J. Polit. Econ.* 107, 731–60.

Clarke, E., (1971). "Multipart Pricing of Public Goods." *Public Choice* 8, 19–33.

Courty, P., and Li, H., (2000). "Sequential Screening." *Rev. Econ. Stud.* 67, 697–717.

Cremer, J., (1995). "Arm's Length Relationships." *Q. J. Econ.* 110, 275–95.

Cremer, J., and McLean, R. P., (1988). "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions." *Econometrica*, 56(6), 1247–57.

Dasgupta, P., Hammond, P., and Maskin, E., (1979). "The Implementation of Social Choice Rules: Some General Results on Incentive Compatibility." *Rev. Econ. Stud.* 66, 185–216.

Deb, R., (2014). "Intertemporal Price Discrimination with Stochastic Values." Discussion paper. University of Toronto.

Deb, R., and Pai, M., (2013). "Ironing in Dynamic Revenue Management: Posted Prices and Biased Auctions." In: *Proceedings of the Symposium on Discrete Algorithms (SODA)* 2013, pp. 620–31.

Deb, R., and Said, M., (2015). "Dynamic Screening with Limited Commitment." *J. Econ. Theory* 159 (Part B), 891–928.

Diamond, P.A., (1998). "Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates." *American Economic Review*, 88(1), 83–95.

Diamond, P.A., and Mirrlees, J.A., (1978). "A Model of Social Insurance with Variable Retirement." *Journal of Public Economics* 10(3), 295–336.

Doepke, M., and Townsend, R., (2004). "Dynamic Mechanism Design with Hidden Income and Hidden Actions," *Journal of Economic Theory* 126 (1), 235–85.

Doval, L., and Ely, J., (2016). "Sequential Information Design," working paper, Northwestern University.

Dustmann, C., and Meghir, K., (2005). "Wages, Experience and Seniority." *Review of Economic Studies*, 72, 77–108.

Edmans, A., and Gabaix, X., (2011). "Tractability in Incentive Contracting." *Rev. Financ. Stud.* 24, 2865–94.

Edmans, A., Gabaix, X., Sadzik, T., and Sannikov, Y., (2012). "Dynamic CEO Compensation." *J. Finance* 67, 1603–47.

Einav, L., Farronato, C., Levin, J., and Sundaresan, N., (2013). "Sales Mechanisms in Online Markets: What Happened to Internet Auctions?" Discussion paper. Stanford University.

Einav, L., Finkelstein, A., and Levin, J., (2010). "Beyond Testing: Empirical Models of Insurance Markets." *Annu. Rev. Econ.* 2, 311–36.

Ely, J., (2015). "Beeps." *American Economic Review*, forthcoming.

Ely, J., Garrett, D., and Hinnosaar, T., (2016). "Overbooking." *Journal of European Economic Association*, forthcoming.

Ely, J., and Szydlowski, M., (2016). "Moving the Goalposts," working paper, Northwestern University.

Eső, P., and Szentes, B., (2007). "Optimal Information Disclosure in Auctions." *Rev. Econ. Stud.* 74, 705–31.

Eső, P., and Szentes, B., (2015). "Dynamic Contracting: An Irrelevance Result." *Theoretical Economics*, forthcoming.

Eső, P., and Szentes, B., (2017). "Dynamic Contracting: An Irrelevance Theorem," *Theoretical Economics* 12, 109–39.

Farhi, E., and Werning, I., (2013). "Insurance and Taxation Over the Life Cycle." *Rev. Econ. Stud.* 80, 596–635.

Fernandes, A., and Phelan, C., (2000). "A Recursive Formulation for Repeated Agency with History Dependence." *J. Econ. Theory* 91, 223–47.

Fershtman, D., and Pavan, A., (2016). "Matching Auctions." Discussion paper. Northwestern University.

Frankel, A., (2012), "Aligned Delegation," *American Economic Review* 2014, 104(1): 66–83.

Frankel, A., (2016), "Discounted Quotas," *Journal of Economic Theory*, forthcoming.

Freixas, X., Guesnerie, R., and Tirole, J., (1985). "Planning under Incomplete Information and the Ratchet Effect." *Rev. Econ. Stud.* 52, 173–91.

Galperti, S., (2015). "Commitment, Flexibility, and Optimal Screening of Time Inconsistency," *Econometrica*, 83(4) 1425–65.

Garrett, D., (2014). "Robustness of Simple Menus of Contracts in Cost-Based Procurement," *Games and Economic Behavior*, 87, September 2014, 631–41.

Garrett, D., (2016a). "Intertemporal Price Discrimination: Dynamic Arrivals and Changing Values." *Am. Econ. Rev.*, forthcoming.

Garrett, D., (2016b). "Dynamic Mechanism Design: Dynamic Arrivals and Changing Values." Discussion paper. Toulouse School of Economics.

Garrett, D., and Pavan, A., (2012). "Managerial Turnover in a Changing World." *J. Polit. Econ.* 120, 879–925.

Garrett, D., and Pavan, A., (2015). "Dynamic Managerial Compensation: A Variational Approach." *J. Econ. Theory* 159 (Part B), 775–818.

Garrett, D., Pavan, A., and Toikka, J., (2016). "Robust Predictions in Dynamic Screening." Discussion paper. Northwestern University, University of Toulouse, and MIT.

Genesove, D., (1993). "Adverse Selection in the Wholesale Used Car Market." *J. Polit. Econ.* 101, 644–65.

Gentzkow, M., and Kamenica, E. (2015). "Information Environments and the Impact of Competition on Information Provision." University of Chicago mimeo.

Gerardi, D., and Maestri, L., (2016). "Dynamic Contracting with Limited Commitment and the Ratchet Effect." Discussion Paper, Collegio Carlo Alberto.

Gershkov, A., and Moldovanu, B., (2009a). "Dynamic Revenue Maximization with Heterogenous Objects: A Mechanism Design Approach." *Am. Econ. J. Microecon.* 2, 98–168.

Gershkov, A., and Moldovanu, B., (2009b). "Learning About the Future and Dynamic Efficiency." *Am. Econ. Rev.* 99, 1576–88.

Gershkov, A., and Moldovanu, B., (2012). "Optimal Search, Learning and Implementation." *J. Econ. Theory* 147, 881–909.

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

Gershkov, A., Moldovanu, B., and Strack, P., (2014). "Revenue Maximizing Mechanisms with Strategic Customers and Unknown Demand: Name-Your-Own-Price." Discussion paper http://ssrn.com/abstract=2527653. Hebrew University and University of Bonn and UC Berkeley

Gershkov, A., and Szentes, B., (2009). "Optimal Voting Scheme with Costly Information Aquisition." *J. Econ. Theory* 144, 1895–1920.

Gibbard, A., (1973). "Manipulation of Voting Schemes." *Econometrica*, 41, 587–601.

Gibbard, A., (1977). "Manipulation of Schemes that Mix Voting with Chance." *Econometrica*, 45, 665–81.

Golosov, M., Troshkin, M. & Tsyvinski, A., (2015). "Redistribution and Social Insurance," *Am. Econ. Rev.*, 106(2), 359–86.

Golosov, M., Troshkin, M., and Tsyvinski, A., (2016). "Optimal Dynamic Taxes." *Am. Econ. Rev.*, 106(2), 359–86.

Golosov, M., Tsyvinski, A., and Werning, I., (2006). "New Dynamic Public Finance: A User's Guide." *NBER Macroeconomic Annual* 2006, MIT Press.

Gomes, R., and Pavan, A., (2016a). "Many-to-Many Matching and Price Discrimination." *Theoretical Economics*, 11, 1005–52.

Gomes, R., and Pavan, A., (2016b). "Matching Plans for Agents with Horizontally and Vertically Differentiated Preferences." Mimeo, Northwestern University.

Green, J., (1987). "'Making Book Against Oneself,' the Independence Axiom, and Nonlinear Utility Theory." *Q. J. Econ.* 98, 785–96.

Green, J., and Laffont, J., (1979). *Incentives in Public Decision Making*. Amsterdam: North-Holland.

Green, J., and Laffont, J., (1987). "Posterior Implementability in a Two-Person Decision Problem." *Econometrica* 55, 69–94.

Groves, T., (1973). "Incentives in Teams." *Econometrica* 41, 617–31.

Grubb, M., and Osborne, M., forthcoming. "Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock." *Am. Econ. Rev.*

Guesnerie, R., and Laffont, J., (1984). "A Complete Solution to a Class of Principal-Agent Problems with an Application to the Control of a Self-Managed Firm." *J. Public Econ.* 25, 329–69.

Gul, F., Sonnenschein, H., and Wilson, R., (1986). "Foundations of Dynamic Monopoly and the Coase Conjecture." *J. Econ. Theory* 39, 155–190.

Guo, Y., and J. Hörner (2016). "Dynamic Contracting Without Money," Discussion paper. Yale University and Northwestern University.

Handel, B., Hendel, I., and Whinston, M., (2015). "Equilibria in Health Exchanges: Adverse Selection vs. Reclassification Risk." *Econometrica*, 83(4), 1261–1313.

Hart, O., and Tirole, J., (1988). "Contract Renogiation and Coasian Dynamics." *Rev. Econ. Stud.* 55, 509–40.

Heathcote, J., Storesletten, K., and Violante, G. L., (2014). "Optimal Tax Progressivity: An Analytical Framework" (Working Paper No. w19899). National Bureau of Economic Research.

Hendel, I., and Lizzeri, A., (2003). The Role of Commitment in Dynamic Contracts: Evidence from Life Insurance." *Q. J. Econ.* 118, 299–327.

Hinnosaar, T., (2016). "Calendar Mechanisms." Discussion paper. Collegio Carlo Alberto.

Holmstrom, B., and Milgrom, P., (1987). "Aggregation and Linearity in the Provision of Intertemporal Incentives." *Econometrica* 55, 303–28.

Hörner, J., and Skrzypacz, A., (2016). "Learning, Experimentation and Information Design." *Adances in Economics and Econometrics. Proceeding of the 2015 World Congress of the Econometric Society*, forthcoming.

Hylland, A., and Zeckhauser, R., (1979). The Efficient Allocation of Individuals to Positions, *Journal of Political Economy*, 87, 293–313.

Jacobson, L. S., LaLonde, R. J., and Sullivan, D. G., (1993). "Earnings Losses of Displaced Workers." *The American Economic Review*, 685–709.

Johnson, T., (2015). "Dynamic Mechanism Without Transfers," working paper, University of Notre Dame.

Kakade, S., Lobel, I., and Nazerzadeh, H., (2013). "Optimal Dynamic Mechanism Design and the Virtual Pivot Mechanism." *Oper. Res.* 61, 837–54.

Kamenica, E., and Gentzkow, M., (2011). "Bayesian Persuasion." *Am. Econ. Rev.* 101(6), 2590–615.

Kapicka, M., (2013a). "Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First Order Approach." *Review of Economic Studies*, 80(3), 1027–54.

Kapicka, M., (2013b). "Optimal Mirrleesean Taxation with Unobservable Human Capital Formation." mimeo, UC Santa Barbara.

Kapicka, M., and Neira, J., (2014). "Optimal Taxation in a Life-Cycle Economy with Endogenous Human Capital Formation." mimeo, UC Santa Barbara.

Kocherlakota, N., (1996). "Implications of Efficient Risk Sharing Without Commitment." *Rev. Econ. Stud.* 63, 595–609.

Kocherlakota, N., (2005). "Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation." *Econometrica* 73, 1587–21.

Kocherlakota, N. R., (2010): *The New Dynamic Public Finance*. Princeton University Press, USA.

Kováč, E., Krähmer, D., and Tatur, T., (2014). "Optimal Stopping in a Principal-Agent Model With Hidden Information and No Monetary Transfers," working paper, University of Bonn.

Krähmer, D., and Strausz, R., (2011). "Optimal Procurement Contracts with Pre-Project Planning." *Rev. Econ. Stud.* 78, 1015–41.

Krähmer, D., and Strausz, R., (2015a). "Optimal Sales Contracts with Withdrawal Rights," *Rev. Econ. Stud.* 82, 762–90.

Krähmer, D., and Strausz, R., (2015b). "Ex Post Information Rents in Sequential Screening," *Games and Economic Behavior* 90, 257–73.

Krähmer, D., and Strausz, R., (2016). "Sequential vs Static Screening – An Equivalence Result." Discussion paper. University of Bonn and Humboldt University of Berlin.

Krause, A., (2009). "Optimal Nonlinear Income Taxation with Learning-by-Doing." *Journal of Public Economics*, 93(9-10), 1098–110.

Krueger, D., and Uhlig, H., (2006). "Competitive Risk Sharing Contracts with One-Sided Commitment." *J. Monet. Econ.* 53, 1661–91.

Kruse, T., and Strack, P., (2015). "Optimal Stopping with Private Information." *J. Econ. Theory* 159 (Part B), 702–27.

Kumar, P., (1985). "Consistent Mechanism Design and the Noisy Revelation Principle." PhD thesis. Stanford University.

Laffont, J., and Tirole, J., (1986). "Using Cost Observation to Regulate Firms." *J. Polit. Econ.* 94, 614–41.

Laffont, J., and Tirole, J., (1988). "The Dynamics of Incentive Contracts." *Econometrica* 59, 1735–54.

Laffont, J., and Tirole, J., (1990). "Adverse Selection and Renegotiation in Procurement." *Rev. Econ. Stud.* 57, 597–625.

Laffont, J., and Tirole, J., (1986). "Using Cost Observation to Regulate Firms." *Journal of Political Economy*, 94, 614–41.

Li, J., Matouschek, N., and Powell, M., (2015). "The Burden of Past Promises," working paper, Northwestern University.

Li, H., and Shi, X., (2015). "Discriminatory Information Disclosure." Discussion paper. University of British Columbia and University of Toronto.

Ligon, E., Thomas, J., and Worrall, T., (2002). "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies." *Rev. Econ. Stud.* 69, 209–44. Discussion paper. University of Rochester.

Liu, H., (2014). "Efficient Dynamic Mechanisms in Environments with Interdependent Valuations," working paper, University of Michigan.

Liu, Q., Mierendorff, K., and Shi, X., (2015). "Auctions with Limited Commitment." Discussion paper. Columbia University and UCL and University of Toronto.

Ljungqvist, L., and Sargent, T., (2004). *Recursive Macroeconomic Theory*, second edition. Cambridge, Massachusetts: MIT Press.

Luz, V., (2015). "Dynamic Competitive Insurance." Discussion paper. University of British Columbia.

Maestri, L., (2016). "Dynamic Contracting Under Adverse Selection and Renegotiation." Discussion paper, Fondacion Gatulo Vergas.

Makris, M., and Pavan, A., (2016). "Incentives for Endogeneous Types: Taxation Under Learning by Doing," Discussion paper. Northwestern University.

Maskin, E., and Riley, J., (1984). "Monopoly with Incomplete Information." *Rand J. Econ.* 15, 171–96.

Maskin, E., and Riley, J., (1989). "Optimal Multi-Unit Auctions." In *The Economics of Missing Markets, Information, and Games*, 312–35. Oxford University Press.

McAfee, R., and Vincent, D., (1997). "Sequentially Optimal Auctions." *Games Econ. Behav.* 18, 24–76.

Miao, J., and Zhang, Y., (2015). "A Duality Approach to Continuous-Time Contracting Problems with Limited Commitment." *J. Econ. Theory* 159 (Part B), 929–88.

Mierendorff, K., (2014). "Optimal Dynamic Mechanism Design with Deadlines." Discussion paper. UCL.

Milgrom, P., and Segal, I., (2002). "Envelope Theorems for Arbitrary Choice Sets." *Econometrica*, 70(2), 583–601.

Miralles, A. (2012). "Cardinal Bayesian Allocation Mechanisms without Transfers," *Journal of Economic Theory*, 147, 179–206.

Miravete, E., (2003). "Choosing the Wrong Calling Plan? Ignorance and Learning." *Am. Econ. Rev.* 93, 297–310.

Mirrlees, J., (1971). "An Exploration in the Theory of Optimum Income Taxation." *Rev. Econ. Stud.* 38, 175–208.

Mueller, C., (2015). "Robust Virtual Implementation under Common Strong Belief in Rationality," *Journal of Economic Theory*, 162, 407–50.

Mussa, M., and Rosen, S., (1978). "Monopoly and Product Quality." *J. Econ. Theory* 18, 301–17.

Myesron, R., (1979). "Incentive Compatibility and the Bargaining Problem," *Econometrica*, 47(1), 61–74.

Myerson, R., (1981). "Optimal Auction Design." *Math. Oper. Res.* 6, 58–73.

Myerson, R. B., (1986). "Multistage Games with Communication," *Econometrica*, 54(2), 323–58.

Myerson, R., and Satterthwaite, M., (1983). "Efficient Mechanisms for Bilateral Trading." *Journal of Economic Theory*, Vol. 29, pp. 265–81.

Nikandrova, A., and Pancs, R., (2015). "An Optimal Auction with Moral Hazard." Mimeo, ITAM.

Pavan, A. (2007). "Long-Term Contracting in a Changing World," Mimeo, Northwestern University.

Pavan, A., Segal, I., and Toikka, J., (2014). "Dynamic Mechanism Design: A Myersonian Approach." *Econometrica* 82, 601–53.

Penta, A., (2015). "Robust Dynamic Implementation," *Journal of Economy Theory,* Vol. 160, 280–316.

Porter, R., (1995). "The Role of Information in U.S. Offshore Oil and Gas Lease Auctions." *Econometrica* 63, 1–27.

Prat, J., and Jovanovic, B., (2014). "Dynamic Contracts when the Agent's Quality is Unknown." *Theoretical Economics* 9, 865–914.

Prendergast, C., (2002). "The Tenuous Trade-Off between Risk and Incentives." *Journal of Political Economy*, 110(5), 1071–102.

Rahman, D., (2011). "Detecting Profitable Deviations." Discussion paper. University of Minnesota.

Rayo, L., and Segal, I., (2010), “Optimal Information Disclosure,” *Journal of Political Economy*, 118(5), 949–87.

Riley, J., and Samuelson, W., (1981). “Optimal Auctions.” *Am. Econ. Rev.* 71, 381–92.

Riordan, M., and Sappington, D., (1987). “Information, Incentives, and Organizational Mode.” *Q. J. Econ.* 102, 243–64.

Rochet, J.-C., (1987). “A Necessary and Sufficient Condition for Rationalizability in a Quasi-Linear Context.” *J. Math. Econ.* 16, 191–200.

Rogerson, W., (1985). “Repeated Moral Hazard,” *Econometrica,* 53(1), 69–76.

Said, M., (2011). “Sequential Auctions with Randomly Arriving Buyers.” *Games Econ. Behav.* 73, 236–43.

Said, M., (2012). “Auctions with Dynamic Populations: Efficiency and Revenue Maximization.” *J. Econ. Theory* 147, 2419–38.

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*. Cambridge: Cambridge University Press, pp. 89–124.

Satterthwaite, M. (1975). “Strategy-Proofness and Arrow’s Condition: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions,” *J. Econ. Theory* 10, 187–217.

Segal, I., (2003). “Optimal Pricing Mechanisms with Unknown Demand.” *American Economic Review*, 93(3): 509–29.

Skreta, V., (2006). “Sequentially Optimal Auctions.” *Rev. Econ. Stud.* 73, 1085–111.

Skreta, V., (2015). “Optimal Auction Design Under Non-Commitment.” *J. Econ. Theory* 159 (Part B), 854–90.

Skrzypacz, A., and Toikka, J., forthcoming. “Mechanisms for Repeated Trade.” *Am. Econ. J. Microecon.*

Stantcheva, S., (2014). “Optimal Income Taxation with Adverse Selection in the Labour Market.” *Rev. Econ. Stud.* 81, 1296–329.

Stantcheva, S., (2016). “Optimal Taxation and Human Capital Policies over the Life Cycle.” *Journal of Political Economy,* forthcoming.

Stokey, N., (1981). “Rational Expectations and Durable Goods Pricing.” *Bell J. Econ.* 12, 112–28.

Strulovici, B., (2016). “Contract Negotiation and the Coase Conjecture.” Discussion paper, Northwestern University.

Strulovici, B., and Szydlowski, M., (2015). “On the Smoothness of Value Functions and the Existence of Optimal Strategies in Diffusion Models.” *J. Econ. Theory* 159 (Part B), 1016–55.

Thomas, J., and Worrall, T., (1988). “Self-Enforcing Wage Contracts.” *Rev. Econ. Stud.* 55, 541–54.

Topel, R., (1991). “Specific Capital, Mobility and Wages: Wages Rise with Seniority.” *Journal of Political Economy*, 99(t), 145–76.

Vickrey, W., (1961). “Counterspeculation, Auctions, and Competitive Sealed Tenders.” *J. Finance* 16, 8–37.

Williams, N., (2011). “Persistent Private Information.” *Econometrica* 79, 1233–74.

Williams, N., (2015). “A Solvable Continuous Time Dynamic Principal-Agent Model.” *J. Econ. Theory* 159 (Part B), 989–1015.

Willis, R., (1986). “Wage Determinants: A Survey and Reinterpretation of Human Capital Earnings Functions,” in *Handbook of Labor Economics*, ed. by O. Ashenfelter and R. Layard, New York: Elsevier Science Publisher, Vol. 1, pp. 525–602.

Wilson, R., (1993). *Nonlinear Pricing*. Oxford: Oxford University Press.