Limited Commitment: Mechanism Design meets Information Design*

Vasiliki Skreta†

November 1, 2025

Abstract

This chapter studies dynamic mechanism design when the principal lacks full commitment and can only offer *spot* mechanisms—contracts governing current interactions without binding future behavior. In such environments, the principal is a strategic player whose re-optimization shapes dynamic incentive constraints and feasible outcomes. The revelation principle of Doval and Skreta (2022) establishes that every equilibrium outcome of a dynamic mechanism-selection game, in which the principal selects at every period a generalized indirect spot mechanisms can be replicated by an equilibrium in which, at every history, the principal offers a *Direct Blackwell Mechanism* (DBM); the agent reports truthfully and participates, and the output message of the mechanism coincides with the principal's belief about the agent. A DBM consists of an information-disclosure rule mapping reported types to posteriors about types, and an allocation rule mapping posteriors to allocations. DBMs tailor allocations to different types while accounting for losses in information rents, formally capturing the agent's privacy and ex post distortion concerns arising from the principal's opportunistic behavior. This revelation principle transforms dynamic mechanism design under limited commitment into constrained information design, providing tractable tools for applications—including regulation, dynamic taxation, vertical relations, and mechanism-design settings with aftermarkets.

Keywords: *mechanism design, limited commitment, revelation principle, direct Blackwell mechanisms, information design, dynamic mechanism design, Markov environments*
JEL classification D84, D86

---

*This chapter heavily draws on joint work with Laura Doval. I am indebted to Laura Doval for countless hours of discussion, detailed feedback on this chapter and for the extraordinary intellectual journey that resulted in the key works overviewed in this chapter. I am privileged to have her as a coauthor and friend. I am grateful to the discussants Tilman Börgers and Roland Strausz whose incisive comments significantly improved this chapter. I also thank Andrea Attar, Francesco Conti, Giovanni De Paola, Dino Gerardi, Alkis Georgiadis-Harris, Nathan Hancart, and Narayana Kocherlakota for feedback and several helpful comments. The usual disclaimer applies.

†University of Texas at Austin, University College London, and CEPR.
E-mail: vskreta@gmail.com

"In practice, regulation is seldom once-and-for-all, but rather an activity that takes place over an extended period. The regulator may then be unable to commit to a regulatory policy over the relevant time span." — Nobel Committee (2014, p. 15), describing Jean Tirole's contributions

1 Introduction

In mechanism design, the designer selects the "rules" of the game—the institution—and the choice is final and committed forever. In multi-period relationships, the principal ex ante commits to a long dynamic mechanism or a sequence of mechanisms for the entire relationship. However, this commitment assumption can be problematic, as institutions, such as contracts, may not remain optimal once information is revealed in future stages. For example, governments may want to reoptimize debt ceilings, firms may want to make personalized offers to customers based on past observed interactions, and employers may want to offer a new contract to a productive employee.

As first discussed in Roberts (1984), time inconsistency in the Mirrleesian optimal nonlinear taxation framework stems from the government's potential exploitation of acquired information. For instance, if agents' skills are binary, the optimal tax scheme distorts low-skilled labor under commitment. However, in the second period, the government may revise taxes once it identifies low-skilled agents.

Despite its shortcomings, commitment is assumed for technical convenience because, under commitment, the mechanism-selection problem can be reduced to a constrained optimization problem thanks to the *revelation principle:* The designer simply selects a mapping from privately-known characteristics to outcomes (a.k.a. *direct revelation mechanism*) subject to constraints that ensure it is optimal for agents to truthfully report their characteristics and participate voluntarily in the mechanism. Unfortunately, this version of the revelation principle fails when the designer cannot fully commit ex-ante. For several decades, researchers faced challenges in settings with imperfect commitment and were forced to impose auxiliary ad hoc assumptions on the available contracts or the length of the contractual horizon or the structure of private information.

In this chapter, we study a principal interacting over T periods with a single privately informed agent whose type may be fully persistent or evolve over time. We focus on *limited commitment*: in each period, the principal offers a spot mechanism to determine current allocations but cannot commit to future mechanisms. This reflects contractual incompleteness: the principal can enforce today's rules, but nothing beyond that. The resulting frictions are central to regulation (Laffont and Tirole, 1988), procurement (Laffont and Tirole, 1993), vertical contracting (Dworczak, 2020), and data privacy settings in which the designer may also act as a data broker (Doval and Skreta, 2025). In all these settings, the inability to commit to future mechanisms introduces informational frictions that shape dynamic incentives and feasible allocations.

Forms of limited commitment While full commitment to all future mechanisms (or a long term contract) represents one benchmark, deviations from it can take several distinct forms, each shaping the set of implementable outcomes in dynamic environments. Table 1 summarizes key scenarios and select contributions from the literature.

Table 1: Forms of imperfect commitment over a horizon of $T \leq \infty$ periods

Commitment type	**Contract structure over T**	**References**
Full commitment	A single contract covering all T periods, fully enforceable and not subject to renegotiation.	Baron and Besanko (1984); Pavan et al. (2014); Bergemann and Välimäki (2019)
Long-term com. w/ renegotiation	A contract specified at $t = 0$ for all T periods, but renegotiable at interim dates with mutual consent.	Dewatripont (1986, 1988); Hart and Tirole (1988); Fudenberg and Tirole (1990)
Short-term commitment	A sequence of contracts of fixed length $1 \leq \tau \leq T$, renewed or renegotiated periodically.	Rey and Salanie (1996); Breig (2022)
Limited commitment	A new one-period contract is offered in each period.	Freixas et al. (1985); Hart and Tirole (1988); Bester and Strausz (2001); Skreta (2006)
No commitment	Even within-period mechanisms may be unenforceable; outcomes must be credible and self-enforcing each period.	McAdams and Schwarz (2007); Akbarpour and Li (2020)

One class of models, pioneered by Dewatripont (1986, 1988), assumes that full-term contracts are specified at the outset but may be renegotiated bilaterally if both parties agree. This framework retains the formal structure of full-term contracts while acknowledging the possibility of ex post inefficiencies that can be alleviated through renegotiation. This is called *long-term commitment under renegotiation.*[1] A second

[1]An important result in that literature is *the renegotiation-proofness principle*, which was estab-

category, called *short-term commitment*, encompasses contracts of intermediate duration—neither purely one-period nor full-term—with renegotiation opportunities arising over their course, as in Rey and Salanie (1996). A third category comprises models in which the principal can only propose spot mechanisms revised each period, capturing complete absence of intertemporal commitment; see Freixas et al. (1985); Laffont and Tirole (1988); Skreta (2006). We refer to this as *limited commitment* or *spot commitment*, the focus of this chapter.[2],[3] Finally, the most stringent form of imperfect commitment arises in models where even the mechanism of today is not fully enforceable—so-called *credible mechanism design*—as studied in McAdams and Schwarz (2007); Vartiainen (2013) and Akbarpour and Li (2020).[4]

The earlier literature suggested that commitment issues can manifest in mild form when parties can enforce long-term contracts but cannot prevent mutually beneficial renegotiation, or in more severe form when only short-term or spot contracts are available (such as when the principal is a government unable to credibly commit against future legislative changes). Under long-term commitment subject to renegotiation, the principal is unable to credibly threaten harsh future treatment of the agent. In contrast, limited commitment prevents commitment to lenient future treatment, generating ratchet effects.

Hart and Tirole (1988) and Breig (2022) compare limited commitment versus long-term contracts under renegotiation in dynamic buyer-seller settings and show that they are outcome equivalent. Perhaps counterintuitively, Breig (2022) shows that spot contracts augmented with random delivery clauses can outperform long-term contracts that are subject to renegotiation. Indeed, in settings with imperfect commitment, the principal is a player in the game, and having more effective deviations available (e.g., deviations to long-term contracts versus only spot contracts) can hurt the principal.

The inability to fully commit harms the principal by creating a wedge between ex ante optimal and sequentially rational outcomes. Information disclosure during contract performance creates renegotiation or reoptimization opportunities that are ex ante detrimental to the principal, generating additional frictions beyond those in second-

lished by Dewatripont (1986) and Hart and Tirole (1988) and states that it is without loss for the principal to offer contracts that are renegotiation-proof.

[2] Laffont and Tirole (1988), Salanié (2005), and Skreta (2006) call this "no commitment." We follow "limited commitment" from Bester and Strausz (2001), subsequently adopted by Doval and Skreta (2022).

[3] A related literature investigates limited commitment constraints within *specific* mechanisms, including dynamic pricing in durable goods markets (Bulow, 1982; Gul et al., 1986; Stokey, 1981) and the evolution of reserve prices in sequential auctions (McAfee and Vincent, 1997; Liu et al., 2019).

[4] There are also some recent papers that endow the principal with long-lived mediators that can store information and release it slowly to the principal (Brzustowski et al., 2023; Lomys and Yamashita, 2022). We discuss those in Section 6.

best contracts under asymmetric information. Salanié (2005) provides an introductory overview of various commitment paradigms, while the seminal Laffont and Tirole (1993) offers a more advanced treatment.

The primary driver of frictions due to limited commitment with time-invariant types is information revealed through contract execution. A central insight emphasizes the value of restricting information revelation through mechanisms featuring partial or complete pooling—indeed, full separation may be impossible (Laffont and Tirole, 1988). A major challenge is that these frictions cause the standard revelation principle to fail. Several approaches have been adopted to sidestep the difficulties stemming from the lack of a canonical class of mechanisms. Freixas et al. (1985) restrict attention to linear incentive schemes, while Laffont and Tirole (1988) analyze particular equilibrium classes without fully characterizing optima. To avoid the ratchet effect, much of the literature in public finance and political economy has focused on time-consistent policies with non-persistent private information (Sleet and Yeltekin, 2006, 2008; Farhi et al., 2012; Golosov and Iovino, 2021). When private information evolves over time (types are imperfectly persistent), the aforementioned frictions weaken (Battaglini, 2007). And when types are drawn independently across periods, the analysis shifts dramatically toward balancing intraperiod risk sharing against intertemporal consumption smoothing.

We now provide a roadmap for the chapter. Section 2 introduces a baseline setup in Section 2.1 and traces the evolution of modeling approaches to limited commitment, beginning with transparent indirect spot mechanisms in Section 2.2—the class studied in several early papers including the seminal contributions of Laffont and Tirole (1988) and of Bester and Strausz (2001)—and showing why the standard revelation principle fails under limited commitment. Section 2.3 discusses the revelation principle of Bester and Strausz (2001), who model limited commitment through reduced-form equilibrium constraints, an approach that proved tractable and enabled groundbreaking results. Section 2.4 introduces noisy indirect spot mechanisms, pioneered by Bester and Strausz (2007), which restore truth-telling by allowing the principal to control information transmission through noisy communication devices. Section 2.5 reviews an alternative outcome-based approach due to Skreta (2006, 2015) that directly characterizes allocations implementable via perfect Bayesian equilibrium (PBE), which extends naturally to multi-agent and continuum type settings.

Section 3 illustrates some key differences between mechanism selection games and the standard mechanism design paradigm. This section acts as a segue to Section 4, which introduces the setting and extensive form studied in Doval and Skreta (2022), the class of generalized indirect mechanisms they define, and proves their revelation principle for limited commitment. Section 4.1 compares the Doval and Skreta (2022) framework with Myerson (1982) and explains why belief recommendations are more suitable than action recommendations. Section 4.2 discusses the extension to continuum type spaces and the issues that arise in establishing revelation principles for mechanism selection games with a continuum of types. Section 4.3 extends the analysis to environments with imperfect persistence, where the agent's type evolves according to a Markov process.

Section 5 applies the revelation principle for limited commitment to the durable goods monopoly problem with binary types (Section 5.1) and continuum types (Section 5.2), demonstrating how limited commitment problems can be recast as a constrained information design problem. Section 5.3 considers abstract settings and uses the constrained information design formulation of the principal's problem under limited commitment to obtain further simplifications. Finally, Section 6 discusses tools and institutions that can mitigate the ratchet effect and outlines directions for future research.

2 Early approaches to limited commitment

This section traces the evolution of modeling approaches to limited commitment in dynamic mechanism design. After describing the setting and introducing key notation in Section 2.1, we begin with *transparent indirect spot mechanisms* in Section 2.2, the initial class used in the literature, and show why the standard revelation principle fails under limited commitment. Section 2.3 introduces the seminal contribution of Bester and Strausz (2001), who establish a revelation principle for reduced-form limited commitment–an approach that proved tractable but left the underlying extensive form implicit. Section 2.4 discusses *noisy indirect spot mechanisms*, pioneered by Bester and Strausz (2007), which allow the principal to control information transmission through noisy communication devices. Section 2.5 reviews an alternative outcome-based approach due to Skreta (2006, 2015). These contributions laid the groundwork for the full game-theoretic treatment and revelation principle for limited commitment due to Doval and Skreta (2022) that we develop in Section 4. Readers interested primarily in recent developments can proceed directly to Section 4 with little loss of continuity.

2.1 A general multi-period principal relationship

We consider a dynamic interaction between a principal (she) and an agent (he) over $T \leq \infty$ periods. To provide a unified treatment, we adopt the notation of Doval and Skreta (2022) throughout, even when discussing results from other papers. The principal holds the bargaining power throughout the relationship. The agent is privately informed about his type $\theta \in \Theta$, which is drawn from a known prior $\mu_1 \in \Delta(\Theta)$. In each period, an allocation $a_t \in A_t$ is chosen. For example, a_t could represent a binary decision, so $a_t \in \{0, 1\}$, or a level of a promotion or a reward (such as a green label), and it can also include tuples such as $a_t = (q, t)$, where q denotes quality or the quantity of a good and t is a transfer, as in the non-linear pricing setting of Mussa and Rosen (1978), or $a_t = (C, t)$, where C is a cost and t a transfer, as in the classical regulation settings surveyed in Laffont and Tirole (1993). We initially assume Θ is at most countable and consider continuum type spaces in Section 4.2. For most of the chapter we assume the agent's type is *fully persistent*—that is, remains constant over time. As foreshadowed in the introduction, this is the case where frictions due to limited commitment are most severe. We discuss evolving private information in Section 4.3.

Outcomes The outcomes of the interaction between the principal and the agent are joint distributions over types and allocations, that is: $\Delta(\Theta \times A^T)$.

Implementable outcomes The set of implementable outcomes depends on the precise details of the extensive form that we specify for the interaction between the principal and the agent and the solution concept (Bayes Nash equilibrium (BNE), Perfect Bayesian equilibrium (PBE), Markov perfect equilibrium (MPE)). The extensive form specifies, as usual, players' choice sets, the precise sequence of moves, and who observes what. The principal's choice set equals the class of mechanisms she can choose from, which to a large extent also determines the agent's choices. The principal's strategy specifies a complete plan of action—namely, the mechanism to be proposed in each period and after every possible history.

Full commitment Myerson (1982) describes the standard mechanism design paradigm as a hybrid of cooperative and noncooperative game theory: the principal controls all possibilities for communication and cooperation and possesses complete bargaining power to commit ex ante to any institutional arrangement, while agents act noncooperatively as utility maximizers who passively accept any equilibrium of the game the principal designs. Under full commitment, the revelation principle (see Myerson, 1986; Sugaya and Wolitzky, 2021; Bergemann and Välimäki, 2019) implies that any feasible joint distribution over $\Delta(\Theta \times A^T)$ can be replicated by an incentive compatible dynamic direct revelation mechanism in which the input messages are type reports and truthful reporting with participation is a best response for the agent.

Key result 1. *Under full commitment, the revelation principle provides two fundamental simplifications: first, it identifies the canonical* message *space—inputs in the* *mechanism are type reports; and second, the agent's canonical* behavior*: full participation and truth-telling.*

Baron and Besanko (1984) study a dynamic relationship between a regulator and a firm under full commitment. Leveraging the revelation principle, they show that the principal's optimal dynamic mechanism specifies the ex-ante optimal allocation in each period.[5] In many settings, however, an allocation that is optimal given the prior may fail to be optimal after information is revealed over time. In other words, an ex-ante optimal mechanism may fail to be sequentially rational (or, equivalently, may be time inconsistent), as was the case in the Mirrleesian taxation setting discussed in the introduction.

Limited commitment Under limited commitment, strategies must be sequentially rational: mechanisms proposed in each period and after every history must form part of a perfect Bayesian equilibrium (PBE). The following section describes key insights from earlier research, including seminal contributions by Laffont and Tirole (1988) and Bester and Strausz (2001, 2007).

2.2 Transparent indirect spot mechanisms

The first-generation papers analyzing dynamic mechanisms under limited commitment endowed the principal with *transparent indirect* spot mechanisms, defined as follows:

Definition 1 (Transparent indirect spot mechanism)**.** *A* transparent indirect spot mechanism *is a pair* $\mathbf{M}_t = (M^{\mathbf{M}_t}, \varphi^{\mathbf{M}_t})$*, where* $M^{\mathbf{M}_t}$ *is the set of input messages and* $\varphi^{\mathbf{M}_t} : M^{\mathbf{M}_t} \to \Delta(A_t)$ *assigns to each input message* $m \in M^{\mathbf{M}_t}$ *a distribution over allocations.*

A *direct revelation mechanism* (DRM) is one where $M^{\mathbf{M}_t} = \Theta$. DRMs are the canonical class under full commitment by the standard revelation principle. DRMs are by definition transparent and importantly allow for randomization over allocations.[6]

Laffont and Tirole (1988); Hart and Tirole (1988); Bester and Strausz (2001, 2000) focused on transparent indirect *deterministic* spot mechanisms.[7] In those papers, the principal observes the agent's message, participation decision, and resulting allocation $a_t(m)$. For example, suppose $T = 2$ and all types participate with probability 1 in the principal's first-period mechanism φ_1. Let $r(m|\theta)$ denote the probability type θ inputs m. Then the principal's posterior is $\mu_2(\theta|m) = \frac{\mu_1(\theta)r(m|\theta)}{\sum_{\bar{\theta}}\mu_1(\bar{\theta})r(m|\bar{\theta})}$. In a PBE, the mechanism

[5] See the elegant and insightful exposition in Chapter 7 of Salanié (2005).

[6] Randomization over allocations is important for the validity of the revelation principle (Myerson, 1982), as underscored by Strausz (2003).

[7] Skreta (2006, 2015) also uses transparent indirect spot mechanisms but allows for randomization.

at $T = 2$ must be a best response given this posterior, but this belief depends on the agent's reporting strategy, which must be optimal given φ_1 and the anticipated mechanism φ_2. It is precisely this interdependence that makes dynamic mechanism design with spot contracts challenging. Laffont and Tirole (1993) write:

> *"The lack of commitment in repeated adverse-selection situations leads to substantial difficulties for contract theory."*

Laffont and Tirole (1988) show in a two-period model with a continuum of types that no separating equilibrium exists—thus, the revelation principle fails because truthful reporting is by definition fully separating. This occurs because the designer cannot credibly commit to not exploit revealed information. Anticipating this, forward-looking agents strategically withhold information—often revealing nothing until the final period—forfeiting information rents if they separate. This "ratchet effect" creates the aforementioned interdependence: future mechanisms condition on information from past ones, so today's mechanism must be optimal given beliefs shaped by both past and future choices. The principal cannot extract rents from high types today without triggering pooling that destroys future screening.[8]

Key result 2 (Failure of the standard revelation principle)**.** *Under limited commitment the standard revelation principle fails.*

2.3 Reduced form limited commitment to transparent indirect mechanisms

Despite the practical relevance of short-term contracting, progress was slow following the publication of Laffont and Tirole (1993)'s book, which superbly overviewed prior work and underscored the difficulties. Progress accelerated after the seminal work of Bester and Strausz (2001), the first paper to provide an abstract analysis of mechanism design with limited commitment and transparent indirect mechanisms and establish a revelation principle.

Bester and Strausz (2001) consider a finite-horizon principal-agent relationship captured in a two-stage reduced form that builds on the classic sender-receiver framework of Crawford and Sobel (1982) by augmenting it with contractual decisions in Stage 1. Stage 2 captures—in a reduced form—limited commitment. In their model, there is a principal and an agent with type $\theta \in \Theta = \{\theta_1, \ldots, \theta_{|\Theta|}\}$, where $2 \le |\Theta| < \infty$

[8]One might hope the taxation principle could simplify analysis by reducing mechanisms to menus of contracts. However, the optimal menu size remains unclear, and agents' randomization over menu items depends on anticipated future menus, reintroducing the same interdependence.

and $\mu_1(\theta_i) > 0$ denotes the prior probability of type θ_i. A sequence of allocations $z = (a_1, a_2) \in Z = A_1 \times A_2$ consists of contractible decisions $a_1 \in A_1$ and non-contractible decisions $a_2 \in A_2$, where A_1 and A_2 are metric spaces. The contractible decision a_1 restricts the principal's feasible choices through a correspondence $F : A_1 \to A_2$; given a_1, the principal can choose $a_2 \in F(a_1)$. Payoffs are $V_i(a_1, a_2)$ for the principal and $U_i(a_1, a_2)$ for type θ_i, both continuous and bounded on Z.

A mechanism specifies a message set M and an allocation rule $\varphi : M \to A_1$, so it is a transparent indirect spot mechanism.[9] The agent's strategy is $q : \Theta \to Q$, where Q is the set of probability measures on M. The principal's strategy is $a_2 : M \to F(a_1(m))$ with posterior beliefs $\mu_2 : M \to \Delta(\Theta)$. An assessment $(q, \mu, a_2, \varphi|M)$ is a perfect Bayesian equilibrium if: (i) the principal's strategy satisfies

$$\sum_{\theta_i} \mu_2(\theta_i|m) V_i(a_1(m), a_2(m)) \geq \sum_{\theta_i} \mu_2(\theta_i|m) V_i(a_1(m), a_2'), \text{ for all } a_2' \in F(a_1(m)), \quad (1)$$

where $\mu_2(\theta_i|m)$ is the posterior that the agent's type is θ_i given message m (which also reveals the same information as conditioning on $a_1(m)$ given that the mechanism is deterministic);[10] (ii) each type's strategy maximizes $\int_M U_i(a_1(m), a_2(m)) dq_i(m)$ given the principal's strategy; and (iii) beliefs satisfy Bayes' rule where applicable: $\mu_2(\theta_i|m) = d(\mu_1(\theta_i) q_i)/d\bar{q}$ with $\bar{q} = \sum_{\theta_i} \mu_1(\theta_i) q_i$. A strategy profile is incentive feasible if it constitutes a PBE given φ, and incentive efficient if no other incentive feasible profile Pareto dominates it from the principal's perspective while holding all types' utilities constant.

Key result 3 (The revelation principle for reduced form limited commitment, and transparent indirect mechanisms Bester and Strausz, 2001)**.** *To sustain incentive efficient payoffs, mechanisms with input messages as type reports ($|M^{\mathbf{M}_t}| = |\Theta|$) are without loss of generality. However, canonical behavior fails: the agent reports truthfully only with positive probability.*

Types typically lie in equilibrium, and the optimal lying strategy remains elusive. Nevertheless, the revelation principle of Bester and Strausz (2001) broke substantial new ground by bounding the input message cardinality. Indeed, several authors rely on their characterization: Bisin and Rampini (2006); Hiriart et al. (2011); Fiocco and Strausz (2015); Beccuti and Möller (2018) study finite-horizon settings with discrete types, while Gerardi and Maestri (2020) consider an infinite-horizon setting, focusing on

[9] As mentioned at the beginning of this section, the class of mechanisms considered in Bester and Strausz (2001) encompasses those studied in most of the limited commitment literature, starting with Laffont and Tirole (1988).

[10] They also assume for simplicity full participation–see footnote 8 in Bester and Strausz (2001)– and do not specify the agent's participation decisions.

equilibria for fully patient players. Most papers in that literature find that the optimum exhibits excessive pooling. However, equilibrium analysis remains challenging. Pinning down how the agent randomizes over messages depends on the designer's future choices, which themselves depend on the designer's posterior beliefs about the agent's type. These beliefs, in turn, depend on the agent's randomization behavior.

While Bester and Strausz (2001) is a foundational contribution, several natural extensions remained open. First, the result that $|M^{\mathbf{M}_t}| = |\Theta|$ is without loss does not extend to multi-agent settings, as Bester and Strausz (2000) demonstrate with an example requiring more inputs than types. Second, the analysis relies on a reduced-form formulation of limited commitment rather than a complete game-theoretic treatment of a fully specified extensive form. The latter is particularly important for studying infinite-horizon settings. Third, and relatedly, their multi-period extension implicitly imposes a Markov structure reflected in Equation 1, invoking the "Principle of Optimality" (Bester and Strausz, 2001, p. 1092). As Ausubel and Deneckere (1989) demonstrate, in infinite-horizon settings non-Markov strategies can be valuable, and such restrictions may not fully characterize the designer's best equilibrium payoff. Fourth, the result holds only for finite type spaces. Finally, they restrict attention to transparent deterministic indirect mechanisms; in subsequent work summarized in the next subsection, they relax the full transparency assumption.

2.4 Noisy indirect spot mechanisms

We now proceed to the next generation, where the focus shifts to spot mechanisms that allow the principal to imperfectly observe the agent's input messages, with the goal of alleviating the effects of contractual incompleteness due to short-term contracting. The key idea, due to Bester and Strausz (2007), is to relax the full transparency assumption and allow for *noisy observability* by augmenting indirect mechanisms with a communication device.

Definition 2 (Noisy indirect spot mechanisms Bester and Strausz, 2007)**.** *A* noisy indirect spot mechanism *consists of an indirect mechanism* $\mathbf{M}_t = (S_t^{\mathbf{M}_t}, \varphi_t)$ *augmented by a* communication device $D = (M_t^{\mathbf{M}_t}, S_t^{\mathbf{M}_t}, \beta)$, *where* $M_t^{\mathbf{M}_t}$ *is the set of input messages,* $S_t^{\mathbf{M}_t}$ *is the set of output messages, and* $\beta : M_t^{\mathbf{M}_t} \to \Delta(S_t^{\mathbf{M}_t})$ *is a mapping from inputs to probability distributions over output messages. The function* φ_t *assigns to each output message* $s \in S_t^{\mathbf{M}_t}$ *a period-*t *allocation; that is,* $\varphi_t : S_t^{\mathbf{M}_t} \to A_t$.

This class of mechanisms encompasses transparent indirect spot mechanisms, which correspond to noisy ones when $M^{\mathbf{M}_t} = S^{\mathbf{M}_t}$ and β is the identity map. The agent chooses reports to maximize expected utility given the stochastic mapping β and anticipating the principal's continuation play. After observing an output message, the principal updates beliefs via Bayes' rule and chooses an action that maximizes expected payoff.

Stage 1 contracts specify output message-contingent allocations; stage-2 allocations are chosen after observing the output message and updating beliefs. While the mapping from output to allocations (φ_t) is deterministic, input messages are randomly assigned to outputs. Equilibrium consists of a reporting strategy, beliefs, and principal strategy forming a PBE. The key insight: noisy communication allows the principal to control information transmission, offering more flexibility than transparent mechanisms and restoring truth-telling. Noise enables both the agent to hide and the designer to tie her hands by learning less.

Key result 4 (Noisy observability restores truth-telling, Bester and Strausz, 2007)**.** *To achieve principal-optimal outcomes in a two-stage reduced form limited commitment and noisy indirect spot mechanism, inputs can be type reports ($M^{\mathbf{M}_t} = \Theta$) with truthful reporting without loss of generality. Under certain conditions, an* upper bound *on output messages is the number of types plus binding monotonicity constraints.*

Beyond two stages, noisy mechanisms become more complex as the agent accumulates private information (the whole past sequence of input messages). Moreover, the cardinality of output messages becomes crucial because current *and* future allocations depend on them. In multi-period settings, multiple continuation equilibria may exist, the designer may randomize over mechanisms, and specifying flexible off-path behavior becomes essential for selecting the principal-optimal outcome. Doval and Skreta (2022) resolve these issues, as we overview in Section 4.

2.5 PBE implementable outcomes approach

An alternative approach sidesteps the question of canonical mechanisms by focusing directly on implementable outcomes. Skreta (2006, 2015) departs from canonical representations and adopts an outcome-based approach inspired by Riley and Zeckhauser (1983). This work studies seller-buyer(s) relationships as in Myerson (1981), but assumes the seller cannot commit to permanently removing the (durable) good from the market if no trade occurs at the initially chosen mechanism. The seller in each period chooses indirect transparent spot mechanisms (allowing for randomization).

The method involves three steps: characterizing the set of equilibrium-feasible outcomes under limited commitment (PBE feasible), optimizing the designer's objective over this set, and constructing an equilibrium that implements the optimal outcome. While this approach extends to multi-agent and non-finite type settings and non-persistent types,[11] the analysis remains intricate. Conceptual challenges persist, including the failure of the taxation principle—since the same allocation today can arise from different input messages revealing different information tomorrow—and the potential need for posteriors with non-convex supports to improve inference. In fact, it is possible to (eventually) fully separate a continuum of types by initially pooling disjoint sets of types. Golosov, Skreta, Tsyvinski, and Wilson (2014) construct such a fully separating equilibrium in a multi-period sender-receiver setting. Moreover, in multi-agent environments, the principal may become privately informed, further complicating the design problem (Skreta, 2015).

Key result 5 (Outcome-based methods, Skreta, 2006, 2015)**.** *Outcome-based methods are applicable to multi-agent settings but remain inherently complex: non-convex supports, endogenously informed principal, new issues with taxation principle.*

To sum up, the first generation of papers on limited commitment reviewed in this section focused on transparent mechanisms (with the exception of Bester and Strausz, 2007, who introduced noisy mechanisms) and yielded important qualitative insights: ratcheting forces may prevent separation, agent mixing is common, and pooling is often optimal. In binary-type settings, both incentive constraints can bind (the "take-the-money-and-run" phenomenon). A recurring theme is how agents mix to hide their private information and avoid future exploitation by the principal. Much of this line of work has been superbly reviewed in Laffont and Tirole (1993); Salanié (2005) and in the survey article by Laffont (1994). Subsequent papers shift the focus to the principal's strategic behavior, which is the key driver of the ratcheting frictions, and put more emphasis on the role of equilibrium multiplicity and how it can be leveraged to tame the principal (Breig, 2022; Doval and Skreta, 2022).

3 Mechanism selection games versus mechanism design

Doval and Skreta (2022) model the dynamic principal–agent relationship under limited commitment as a mechanism selection *game*—an extensive-form game in which the principal selects mechanisms. Their revelation principle replicates *all* equilibrium outcomes rather than only incentive-efficient ones. Beyond these crucial differences, they depart from Bester and Strausz (2001, 2007) by endowing the principal with a larger class of mechanisms. Before delving into their setting and results, we illustrate some

[11] The outcome-based method is used, for example, by Deb and Said (2015), who study sequential screening under short-term commitment.

important differences between mechanism selection games and the standard mechanism design paradigm.

In standard mechanism design, the principal has complete bargaining power to commit ex ante to any institutional arrangement maximizing her objective. By contrast, Doval and Skreta (2022) study mechanism selection games. Because the principal is now a player rather than a designer who commits ex ante to a set of rules governing all outcomes, three key distinctions arise relative to both the standard mechanism design paradigm and standard dynamic *decision* problems (which by definition do not involve strategic interaction). First, in mechanism selection games, as nicely illustrated by Breig (2022), the principal may benefit from selecting from a smaller set of admissible mechanisms: appropriately restricting choices can tie the principal's hands ex post and lead to better ex ante outcomes. Second, the principal may prefer not to optimize in every period; the principal-optimal PBE may involve continuation play that does *not* maximize the principal's payoff at that stage of the game. This is why replicating all outcomes is important even if we are only interested in principal-optimal ones. Third, the principal may prefer less information to more, as some noise or opaqueness can alleviate the effects of future discretion and re-optimization (cf. Crémer, 1995).[12] These distinctions reveal that dynamic mechanism selection games under limited commitment bear similarities with dynamic decision problems involving time-inconsistent preferences (Strotz, 1955).

Let us elaborate. In a decision problem, enlarging the set of feasible choices is always weakly beneficial; in a strategic environment, however, this need not hold. A larger set of mechanisms grants the principal flexibility to choose the best option today, but also introduces additional opportunities to deviate in the future. Restricting the set of mechanisms can therefore serve as a form of commitment and bears the flavor of a *Ulysses pact*—a self-imposed constraint that binds one's future self to a chosen course of action.[13] This logic illustrates how the principal may benefit from credible self-restraint: the agent is more willing to cooperate today when anticipating that the principal will not exploit the relationship tomorrow. Moreover, in standard dynamic decision problems, the principle of optimality implies that choices must maximize the decision maker's continuation payoff at every period. This is not the case in dynamic mechanism design problems because the information revealed over time *changes* the

[12]The strategic use of information structures as a commitment device to ameliorate the effects of limited commitment was first explored by Crémer (1995) and Dewatripont and Maskin (1995). Subsequently, Bester and Strausz (2007) formalized this idea by endowing the principal with a noisy communication device, and Doval and Skreta (2022) further generalized the approach by introducing mechanisms that jointly control current allocations and information released in subsequent periods.

[13]The term refers to the Greek hero Odysseus (Ulysses), who had himself bound to the mast of his ship to resist the sirens' song.

principal's objective and what is ex post optimal may be ex ante detrimental. Relatedly, it may be necessary to characterize even pessimal continuation equilibria in order to identify the principal-optimal PBE. Both phenomena stand in contrast to standard mechanism design, where a larger design space is unambiguously advantageous and *design* is synonymous with *optimization* and more information is always better.

We illustrate these points through a simple—albeit somewhat contrived—example, chosen to convey the underlying ideas in the most transparent possible setting.

Two-period example. Consider a simple two-period setting. The agent's type θ is binary, $\Theta = \{1, 3\}$, with a uniform prior. Period 1 allocations a_1 consist of a probability of trade q and a transfer x. Rejection corresponds to $q = 0$ and $x = 0$. The period 2 allocation $a_2 \in A_2 = [0, 1]$, with rejection corresponding to $a_2 = 0$. The agent's period 1 payoff is

$$U(a_1, a_2, \theta) = u_1(a_1, a_2, \theta) + u_2(a_2, \theta)$$

where

$$u_1(a_1, a_2, \theta) = \begin{cases} q\theta - x & \text{if } a_1 = (q, x); a_2 = 0, \\ 0 & \text{if } a_1 = (q, x); a_2 \neq 0 \end{cases}$$

and $u_2(a_2, \theta) = 0$ for all a_2, implying that in period 2 the agent is indifferent among all allocations, including rejection. The principal's payoff is additively separable

$$W(a_1, a_2, \theta) = w_1(a_1, \theta) + w_2(a_2, \theta)$$

with $w_1(a_1, \theta) = x$, i.e., the revenue from the period 1 transfer. In period 2, the principal's payoff is $w_2(a_2, \theta) = a_2 \frac{\theta}{K}$ for $K \geq 3$.

A story behind the aforementioned (admittedly extreme) payoff structure is as follows: Consider a gym that offers memberships. In period 1, consumers decide whether to join (with probability q) at price x, anticipating future treatment. The consumer's type θ captures valuation for the gym's services. In period 2, the gym sets intensity of follow-up marketing a_2 through aggressive upselling, pushy marketing, or intrusive contact, generating revenue $w_2(a_2, \theta)$ from supplements or upgrades. Since consumers have already joined and cannot easily exit, they are stuck regardless of harassment level, so $u_2(a_2, \theta) = 0$ for all a_2. Prospective members, however, anticipate such treatment and will join only if they expect to be left alone ($a_2 = 0$).

Commitment optimum Under full commitment, the principal-optimal direct revelation mechanism is:

$$q(1) = 0, \quad x(1) = 0, \quad a_2(1) = 0; \qquad q(3) = 1, \quad x(3) = 3, \quad a_2(3) = 0.$$

This amounts to posting a price of 3 that only the high type accepts, yielding expected payoff $0.5 \times 3 = 1.5$, while committing to zero harassment ($a_2(\theta) = 0, \forall \theta$).

Naive sequential rationality constraint Suppose $T = 2$. Naive sequential rationality requires that the period 2 allocation maximizes the principal's payoff at the beginning of period 2. In the context of our example, this constraint requires

$$a_2 \in \arg\max_{a'_2 \in [0,1]} \mathbb{E}_{\mu_2} w_2(a'_2, \theta) = \mathbb{E}_{\mu_2} a'_2 \frac{\theta}{K}. \tag{N-SR}$$

The unique maximizer is $a_2 = 1$. Since $a_2 = 1 \neq 0$, the agent's period 1 payoff becomes zero (the agent obtains positive utility only when $a_2 = 0$). Consequently, the principal's period 1 payoff is at most zero, leaving only the period 2 payoff, which is $\mathbb{E}[w_2(1, \theta)] = \mathbb{E}[\theta]/K \leq 1$.

Continuation equilibria in period 2 The game-theoretic analysis of Doval and Skreta (2022), who consider a possibly infinitely long interaction, seeks to replicate all equilibrium and continuation-equilibrium outcomes. Their approach requires that mechanisms at every t and at every history are part of a PBE. Let us illustrate the implications in this example. Because $u_2(a_2, \theta) = 0$ for all a_2, the agent is indifferent between accepting and rejecting any period 2 mechanism. Hence, any $a_2 \in [0, 1]$ can arise as part of a continuation-equilibrium outcome. Moreover, rejection of any period-2 mechanism (which yields $a_2 = 0$) is always a best response for the agent. The principal can therefore implement the commitment optimum.

Naive sequential rationality vs. continuation equilibrium approach For any period 1 mechanism, the naive sequential rationality constraint (N-SR) forces $a_2 = 1$, limiting expected payoff to at most 1. In contrast, the continuation-equilibrium approach achieves payoff 1.5, coinciding with the commitment optimum. Thus, (N-SR) is *not* equivalent to continuation-equilibrium feasibility and imposes an ad hoc Markov restriction with loss of *optimality.*[14],[15] This example, though deliberately stylized, illustrates the distinction in the simplest two-period setting.

Returning to the gym story: under (N-SR), the gym sets maximal harassment $a_2 = 1$ in period 2, which either deters enrollment in period 1 or, even if the agent joins in period 1, she is willing to pay at most 0. Under the continuation equilibrium approach, the gym can credibly commit—through equilibrium selection—to $a_2 = 0$, exploiting the multiplicity of continuation equilibria. This allows surplus extraction upfront via membership fees while maintaining a reputation for restraint.

It is easy to modify the example to show that restricting the choices for the principal can help: Trivially, if we restrict $A_2 = \{0\}$, the principal achieves the full commitment optimum even under the constraint (N-SR).

4 The revelation principle under limited commitment

The material in this section builds on and draws heavily from previous joint work with Laura Doval.[16] Doval and Skreta (2022) consider the general setting laid out in Section 2.1. It is worthwhile noting that they do not impose time separability or discounting—these are special cases of their setting. Moreover, their modeling of allocations allows for technological evolution of feasible allocations. In particular, at any period $t \geq 1$, allocations arise from a set $\mathcal{A}_t(a^{t-1})$, where $a^{t-1} = (a_1, \ldots, a_{t-1})$ denotes the history of past allocations, and the outside option, denoted a^*, is always available. Letting $a^T \in A^T$ denote the full allocation path and $\theta \in \Theta$ the agent's type, payoffs are $W(a^T, \theta)$ and $U(a^T, \theta)$ for the principal and agent, respectively.

Doval and Skreta (2022) analyze an extensive form game in which the principal in each stage selects among *generalized spot indirect mechanisms* which they define as follows:

Definition 3 (Generalized spot indirect mechanism Doval and Skreta, 2022)**.** *A generalized indirect spot mechanism is a triple* $\mathbf{M}_t = (M^{\mathbf{M}_t}, S^{\mathbf{M}_t}, \varphi^{\mathbf{M}_t})$ *where*

$$\varphi^{\mathbf{M}_t} : M^{\mathbf{M}_t} \to \Delta(S^{\mathbf{M}_t} \times A_t),$$

[14] Ausubel and Deneckere (1989) first demonstrated this point within an infinitely-long price-posting game.

[15] Naive sequential rationality is without loss of optimality when there is a unique continuation equilibrium or when there is a maximizer that coincides with the ex ante desired continuation equilibrium.

[16] The exposition follows our earlier papers Doval and Skreta (2022, 2024b) closely, though the treatment here is necessarily more cursory; readers seeking complete details are referred to those sources.

assigns to each input message $m \in M^{\mathbf{M}_t}$ *a probability distribution over output messages* $S^{\mathbf{M}_t}$ *and allocations* A_t*.*

Generalized spot mechanisms jointly randomize over output messages and allocations. To make the principal's strategy well-defined, they equip the principal with a collection $\{(M_i, S_i)\}_{i \in \mathcal{I}}$ of input and output message spaces, satisfying the following conditions: (i) each M_i and S_i is a Polish space, (ii) the cardinality of the type space satisfies $|\Theta| \leq |M_i|$ and each M_i is at most countable, and (iii) the cardinality of the posterior space satisfies $|\Delta(\Theta)| \leq |S_i|$. Additionally, they assume that the pair $(\Theta, \Delta(\Theta))$ is included in this collection. Let $\mathcal{M}_{\mathcal{I}}$ denote the set of all mechanisms whose message spaces belong to the collection $\{(M_i, S_i)\}_{i \in \mathcal{I}}$. That is, $\mathcal{M}_{\mathcal{I}} = \bigcup_{i,j \in \mathcal{I}} \{\varphi : M_i \to \Delta(S_j \times A) \mid \varphi \text{ is measurable}\}$. A mechanism-selection game is indexed by these families and is denoted $G_{\mathcal{I}}$.

Timing of mechanism selection-game In each period t,

- t.0 The principal and the agent observe the realization of a public randomization device, $\omega \sim U[0,1]$.
- t.1 The principal chooses a generalized spot indirect mechanism $\mathbf{M}_t = (M^{\mathbf{M}_t}, S^{\mathbf{M}_t}, \varphi^{\mathbf{M}_t})$.
- t.2 Upon observing the mechanism, the agent decides whether to:
    - a. **Accept:** The agent then *privately* reports a message $m \in M$. Subsequently, a pair (s, a) is drawn from $\varphi^{\mathbf{M}_t}(\cdot \mid m)$ and *publicly* observed.[17] The game then proceeds to period $t+1$.
    - b. **Reject:** The agent receives the outside option a^*. The game then proceeds to period $t+1$.

The first key distinction from earlier work is that Doval and Skreta (2022) characterize *all* equilibrium outcomes of a dynamic mechanism selection game. This broader characterization is necessary even when focusing on the principal-optimal equilibrium as already mentioned in Section 2.3 and illustrated by the example in Section 3. A second, key distinction from earlier work lies in the class of mechanisms available to the principal. In the setting in Doval and Skreta (2022), a mechanism, as a function of input messages, *jointly* determines a lottery on the allocations and the output

[17]Output signals are meant to capture the noise the mechanism adds to an agent's report—the transparency of an institution. The modeling of output signals as public ensures that the transparency of the mechanism is commonly known.

messages. Finally, an important feature of the extensive form specified by Doval and Skreta (2022) is the public randomization device.

Doval and Skreta, 2022 focus on assessments $(\sigma_P, \sigma_A, \mu)$ that constitute a PBE of the mechanism selection game(s). At a PBE, the principal's and agent's strategies are sequentially rational at each history, and beliefs are determined via Bayes' rule whenever possible (see Doval and Skreta, 2022 for precise definitions). The principal's histories coincide with *public histories*

$$h^t = (\omega_1, \mathbf{M}_1, \pi_1, s_1, a_1, \cdot, \omega_{t-1}, \mathbf{M}_{t-1}, \pi_{t-1}, s_{t-1}, a_{t-1}, \omega_t)$$

that contain the outcome of the public randomization device, mechanisms, the agent's participation choice (π_t), and the output and allocation drawn from the mechanism. The agent's private histories include the past input messages chosen by the agent:

$$h_A^t = (\omega_1, \mathbf{M}_1, \pi_1, m_1, s_1, a_1, \cdot, \omega_{t-1}, \mathbf{M}_{t-1}, \pi_{t-1}, m_{t-1}, s_{t-1}, a_{t-1}, \omega_t) \in H_A^t(h^t),$$

where $H_A^t(h^t)$ denotes the set of agent private histories possible given a public history h^t. Consequently, the agent accumulates private information as the game proceeds.

The principal's strategy $\sigma_{Pt}(h^t)$ describes the principal's (possibly random) choice of mechanism at time t, which at a PBE must be optimal given beliefs about the agent's private information at period t, denoted by $\mu_t(h^t) \in \Delta(\Theta \times H_A^t(h^t))$. Let $\sigma_{At}(\theta, h_A^t, \mathbf{M}_t) \equiv \big(\pi_t(\theta, h_A^t, \mathbf{M}_t), r_t(\theta, h_A^t, \mathbf{M}_t)\big)$ denote the agent's participation decision and reporting strategy when the principal offers $\mathbf{M}_t$. As we see in what follows, the revelation principle in Doval and Skreta (2022) establishes that it is without loss of generality to focus on *public* PBEs which are PBEs in which both players condition play on the public history.

Equilibrium outcomes The prior μ_1 and the assessment $(\sigma_P, \sigma_A, \mu)$ induce a distribution over the terminal nodes of the game, which we project to $\Theta \times A^T$.

To state the revelation principle in Doval and Skreta (2022), which establishes a canonical class of mechanisms and canonical behavior that suffice to replicate *all* equilibrium outcomes of $G_{\mathcal{I}}$ which we denote by $\mathcal{O}_{\mathcal{I}}^*$, let us define (i) the canonical game, (ii) canonical mechanisms, and (iii) canonical assessments (canonical equilibrium behavior).

The *canonical game* is the mechanism-selection game in which $\mathcal{I} = \{(\Theta, \Delta(\Theta))\}$. Denote the set of equilibrium outcomes of the canonical game by $\mathcal{O}^*$. The *canonical mechanisms* for the Doval and Skreta (2022) setting are a class they define as *Direct-*

Blackwell Mechanisms. We proceed with their definition and a comparison with the standard direct revelation mechanisms:

Definition 4 (Direct-Blackwell Mechanisms Doval and Skreta, 2022)**.** *A mechanism* $\varphi : \Theta \to \Delta(\Delta(\Theta) \times A)$ *is a* direct-Blackwell mechanism (DBM) *if it can be written as*

$$\varphi(\mu, a \mid \theta) = \beta(\mu \mid \theta)\alpha(a \mid \mu),$$

where

$$\beta : \Theta \to \Delta(\Delta(\Theta)), \quad \alpha : \Delta(\Theta) \to \Delta(A)$$

are the mechanism's disclosure *and* allocation *rules.*

Let $\mathcal{M}_C$ (where C stands for canonical) denote the set of DBMs.

Direct-Blackwell mechanisms separate information design (the mapping β) from allocation design (α): β is a statistical experiment encoding the information the principal learns about θ, and conditional on this information, the allocation is drawn independently of θ. Thus, a DBM encodes both the allocation rules and the information the designer obtains from the interaction with the agent.

Remark 1 (Direct Revelation Mechanisms are Direct-Blackwell Mechanisms)**.** *A direct revelation mechanism (DRM):*

$$g : \Theta \to \Delta(A)$$

corresponds to the following DBM:

$$\beta(\cdot \mid \theta) = \delta_\theta, \quad \alpha(\cdot \mid \delta_\theta) = g(\theta).$$

A DRM may reveal too much—for instance, consider two types θ and θ' such that $g(\theta) = g(\theta')$: a DRM necessarily reveals them as distinct, even though they are pooled to the same allocation. This, however, is generally not a concern: under commitment, the agent's (continuation) payoffs do not depend on the beliefs the principal holds about her type after the agent reports, because the mapping from reports to allocations is already set in stone.

Canonical assessments A canonical assessment defines straightforward behavior for both the principal and agent. It is a public PBE in which the principal selects DBMs at every history (on and off the path). When faced with the principal's equilibrium mechanism choice, the agent participates and reports honestly. Finally, when the mechanism generates output $\mu \in \Delta(\Theta)$, this μ represents the principal's revised beliefs regarding the agent's type. In formal terms:

Definition 5 (Canonical assessments Doval and Skreta, 2022)**.** *An assessment*

$(\sigma_P, \sigma_A, \mu)$ *of mechanism-selection game* $G_{\mathcal{I}}$ *is canonical if the following holds for all* $t \geq 1$ *and all public histories* h^t*:*

1. *The principal offers canonical mechanisms, that is,* $\sigma_{Pt}(h^t)(\mathcal{M}_C) = 1$.
2. *For all mechanisms* $\mathbf{M}_t$ *in the support of the principal's strategy at* h^t*,*
   - (a) *All types* θ *in the support of the principal's beliefs in period* t *participate,* $\mu_t(h^t)$*,* $\pi_t(\theta, h_A^t, \mathbf{M}_t) = 1$*,*
   - (b) *All types* θ *in* Θ*, report truthfully* $r_t(\theta, h_A^t, \mathbf{M}_t) = \delta_\theta$*, and*
   - (c) *The mechanism's output belief* μ *coincides with the principal's updated belief about the agent's type. Formally, for* $h^{t+1} = (h^t, \mathbf{M}_t, 1, \mu, \cdot)$*, the marginal of* $\mu_{t+1}(h^{t+1})$ *on* Θ*,* $\mu_{t+1\Theta}(h^{t+1})$*, coincides with* μ*.*[18]
3. *The agent's strategy depends only on his private type and the public history.*

Denote the set of equilibrium outcomes arising at canonical assessments of the canonical game by $\mathcal{O}^C$. Note that canonical assessments are public PBE's.

Theorem 1 (Revelation principle Doval and Skreta, 2022)**.** *For any PBE outcome of any mechanism-selection game* $G_{\mathcal{I}}$*, an outcome-equivalent canonical PBE of the canonical game exists. That is,*

$$\bigcup_{\mathcal{I}} \mathcal{O}^*_{\mathcal{I}} = \mathcal{O}^* = \mathcal{O}^C.$$

Theorem 1 establishes that one can, without loss of generality, restrict attention to mechanisms in which communication is direct: that is, the message space coincides with the set of types, $M = \Theta$. In equilibrium, communication is truthful. Moreover, output messages represent beliefs, so that $S = \Delta(\Theta)$. In equilibrium, these output messages coincide with the principal's equilibrium beliefs. Canonical play reduces the effects of the agent's strategic behavior to participation, truth-telling, and Bayes plausibility constraints. And because canonical PBEs are public, recursivity follows, which is important for tractability.

Theorem 1 provides a key tool for characterizing optimal mechanisms in the presence of limited commitment. A cornerstone of the result is the idea that a mechanism should encode not only the current allocation but also the information the designer obtains from the interaction. Therefore, how much the designer learns–the key distinct tension in design with limited commitment–becomes an explicit part of the design.

[18]The 1 in $(h^t, \mathbf{M}_t, 1, \mu, \cdot)$ denotes the agent's decision to participate in mechanism $\mathbf{M}_t$.

Comparison with standard revelation principle Theorem 1 proves that canonical equilibria of the canonical game–the game when the designer selects mechanisms with inputs equal to type reports and outputs equal to beliefs about the agent's type replicate any equilibrium outcome of any mechanism-selection game, analogous to the standard revelation principle. However, unlike the standard revelation principle, the set of distributions over $\Theta \times A^T$ that constitute equilibrium outcomes depends on both the principal's and the agent's payoffs. As a result, one cannot treat the principal as maximizing an arbitrary objective function over $\mathcal{O}^C$–the principal is a player in the game and her payoffs determine $\mathcal{O}^C$.

Proof overview Recall the usual steps: a direct revelation mechanism results from composing the principal's indirect mechanism with the agent's participation and reporting strategy. In the Doval and Skreta (2022) case, this composition takes the following form. Suppose the principal offers a mechanism $\varphi_t : M \to \Delta(S \times A)$; then the induced mapping is

$$\varphi'_t(s_t, a_t \mid \theta, h^t_A) = (1-\pi_t(\theta, h^t_A, \mathbf{M}_t))\, \mathbb{1}[(s_t, a_t) = (\varnothing, a^*)] + \pi_t(\theta, h^t_A, \mathbf{M}_t) \sum_{m \in M} r_t(\theta, h^t_A, \mathbf{M}_t)(m)\, \varphi_t(s_t, a_t \mid m),$$

which defines a mapping φ' from $\Theta \times H^t_A(h^t)$ to $\Delta\left(S \times A \cup \{(\varnothing, a^*)\}\right)$, where $(\varnothing, a^*)$ is the output signal and the allocation resulting when the agent does not participate $(\pi_t(\theta, h^t_A, \mathbf{M}_t) = 0)$. See Figure 1.

Input messages Let us begin by identifying the appropriate domain of input messages. At first glance, one might consider $M = \Theta \times H^t_A(h^t)$, since the agent's full private history may be relevant (recall that the agent accumulates private information over time). However, Doval and Skreta (2022) ultimately show that messages of the form $M = \Theta$ suffice, relying on an auxiliary result (Proposition B.1 in that paper) that establishes the following: for every perfect Bayesian equilibrium (PBE) assessment of $G_{\mathcal{I}}$, there exists an outcome-equivalent PBE assessment such that the agent's strategy depends only on θ and the public history. The reason is that the part of the agent's private information beyond θ consists of past input messages which, conditional on the past realized allocations and output messages, are payoff-irrelevant. Hence, since the private history (beyond θ) is payoff-irrelevant once we condition on public information, a type θ with private history h^t_A must get the same payoff as type θ with private history $\overline{h}^t_A$. Therefore, we can have the agent's strategy depend only on θ and public history, not on the private history (Figure 1 and Figure 2). In a sense, conditioning on the private history is analogous to conditioning on a public randomization device.

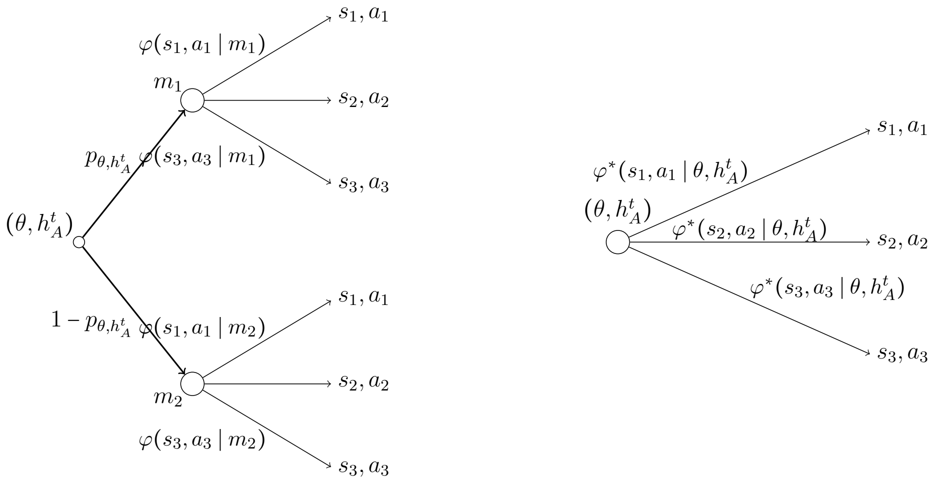

Figure 1: Input message: full private information (θ, h_A^t)

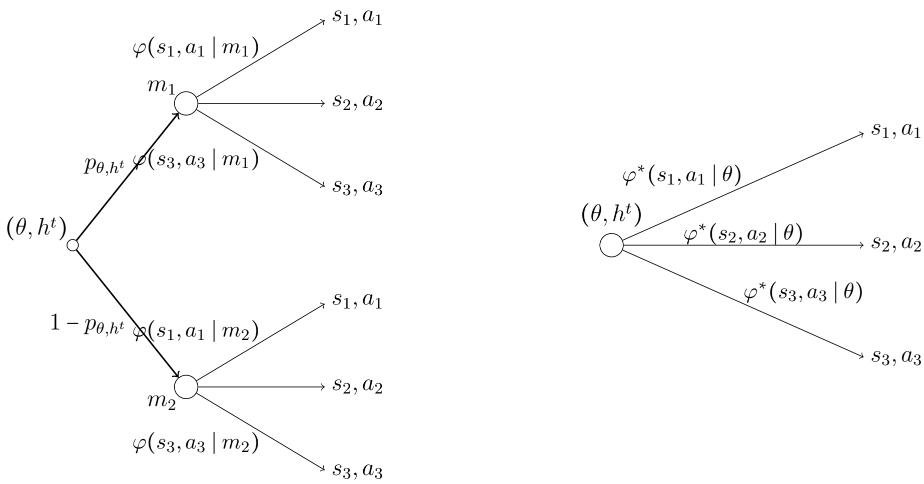

Figure 2: Input message as θ suffices

Output messages We now proceed to explain why it is sufficient to take $S = \Delta(\Theta)$. To do so, we show that any pair (s, a)—where s is a signal and a is an allocation—can be equivalently represented by a pair (μ, a), where μ is the posterior belief about θ derived from s and the prior. One difficulty arises from the fact that the principal can use the output signal s not only to choose allocations but also to coordinate future play. In particular, different signals associated with the same posterior belief about θ may lead to different continuation play. By using public randomization, the principal can simulate this coordination while conditioning only on the publicly observable posterior belief μ (Figure 5). Second, two different output messages inducing the same posterior may generate different distributions over allocations. This randomness can be built into the allocation rule (Figure 4).

Importantly, since we are looking at signals that induce the same belief μ, Bayes' rule ensures the mechanism can be decomposed into two parts: (1) a mapping β from θ to a distribution over posterior beliefs (i.e., a Blackwell experiment), and (2) a mapping from μ to allocations. In other words, in a direct-Blackwell mechanism, the allocation is drawn independently of θ conditional on μ (Figure 4). Thus, posterior beliefs act as sufficient statistics: once μ is fixed, the allocation and the continuation play are independent of θ (Figure 5).

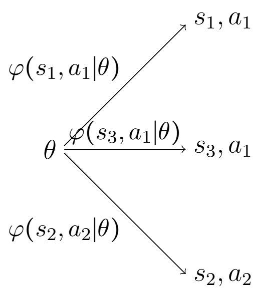

Figure 3: A generalized spot mechanism

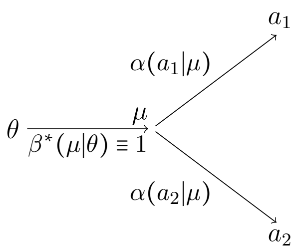

Figure 4: Direct Blackwell mechanism: disclosure and allocation rule

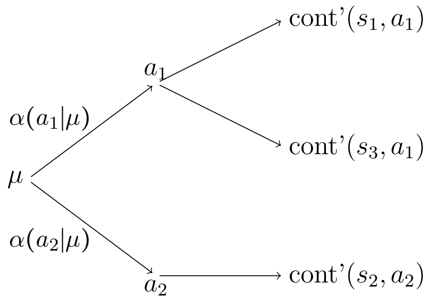

Figure 5: Continuation play coordinated via randomization device

Participation with probability one A central force that gives rise to Direct-Blackwell mechanisms is Bayes' rule which underpins the conditional independence between types and allocations conditional on beliefs, as well as between types and continuation outcomes conditional on both beliefs and observed allocations. This structure arises because the principal's equilibrium beliefs themselves are updated via Bayes' rule. One implication is that output messages that are used only by types with zero posterior probability are eliminated from the mechanism. It follows that types with zero posterior probability—those not in the support of the principal's updated beliefs—may not want to participate anymore. Instead, the mechanism ensures participation only for those types that lie within the support of the principal's beliefs. This is an implication of the consolidation of output messages to beliefs.

4.1 Discussion: beliefs as canonical output messages

This subsection goes deeper into why beliefs suffice as canonical output messages by further clarifying the role of public randomization in allowing canonical mechanisms to replicate all equilibrium outcomes despite using the coarse language of beliefs. It also compares the Doval and Skreta (2022) framework with Myerson (1982)'s general principal-agent model and explains why belief recommendations are more suitable than action recommendations as output messages in mechanism selection games. Readers may skip this discussion without loss of continuity.

In a sense, the "language" of DBMs is very coarse: First, output messages are beliefs, whereas output messages can encode future and past play among many other possibilities. Second, *many* indirect generalized mechanisms can lead to the *same* DBM (recall the composition argument and Figure 1). However, the public randomization device allows us to subsume these additional possibilities for coordination as we proceed to explain.

Public randomization The randomization device serves two coordination roles. First, it allows the principal to coordinate continuation play when multiple continuation equilibria are consistent with the same posterior belief—a role that output messages can play in the mechanism-selection game but not in canonical equilibria, where outputs must coincide with posteriors. Second, it encodes which mechanism was offered when different indirect mechanisms lead to the same direct-Blackwell mechanism, thereby preserving the distinct continuation play associated with each original mechanism. Together, these roles ensure that canonical mechanisms, despite using a coarser language than indirect mechanisms, can replicate the full outcome distribution from the mechanism-selection game. In a sense, it allows us to "distill" all randomness that depends on the agent's type onto the output message s, which is a posterior, and everything orthogonal to the agent's type onto the public randomization device. It also allows for tractable purification of the principal's strategy.

Beliefs versus action recommendations The formulation of generalized spot indirect mechanisms is inspired by Myerson (1982), who studies principal-agent problems involving both enforceable and non-enforceable decisions. In that framework, enforceable decisions correspond to current allocations A_t, while non-enforceable ones correspond to what the $t+1$ principal chooses which are guided by the current mechanism's output message. As Myerson (1982) argues, one can restrict input messages to type reports and output messages to *private* action recommendations without loss of generality. We explain why setting $S^{\mathbf{M}_t}$ to be action recommendations is problematic here.

In Myerson (1982), mechanisms randomize over allocations and output messages privately revealed to each agent. These messages guide agents' actions not directly controlled by the principal. With limited commitment, it is as if we have different principals in each period, and the period-1 principal cannot control her period-2 incarnation. We could attempt to map a two-period mechanism-selection game under limited commitment into Myerson (1982)'s setting by interpreting the period-2 principal as an agent who observes a signal and selects an action (recall our parallel to a receiver in Section 2.4). However, the period-2 principal also observes the period-1 allocation (the contractible decision in Myerson, 1982). More importantly in the Doval and Skreta (2022) extensive form, the period-2 principal is not constrained by a fixed action set but instead endogenously designs *mechanisms*, which creates problems if one equates output messages to action recommendations.

Referring to the action set as generalized spot indirect mechanisms becomes self-referential and conceptually problematic: the output message $S^{\mathbf{M}_t}$ (which we call the output message and which Myerson equates to action recommendations) would itself be a mechanism, that is, a triple $(M^{\mathbf{M}_t}, S^{\mathbf{M}_t}, \varphi : M^{\mathbf{M}_t} \to A \times S^{\mathbf{M}_t})$. Moreover, it is unclear how to handle multiple continuation equilibria for a given mechanism. Coupling a mechanism with different equilibrium play yields different continuation outcomes, and how we handle this multiplicity is crucial for optimality, as we saw in Section 3. One might try to recommend continuation outcomes directly, but this approach fails because continuation outcomes depend on what mechanisms the principal can propose, and those mechanisms themselves have output messages that need specification.

Even in two-period environments—where the standard revelation principle pins down the canonical class for $T = 2$—identifying which DRM to recommend remains nontrivial. More importantly, even a direct mechanism has multiple equilibria, and it may not be enough to suggest a mechanism. As Skreta (2006, 2015) shows, this difficulty arises because the second-period mechanism's optimality depends on posterior beliefs shaped jointly by the first-period mechanism and agents' anticipations of the second-period mechanism (recall also the discussion in Section 2.2). The canonical language of beliefs in Doval and Skreta (2022) resolves these difficulties by separating the design of information and allocations and provides the tractable formulation that allows one to use information design techniques as we overview in Section 5.

4.2 Continuum of types

In this subsection, we overview mechanism-selection games where agent types are drawn from a continuum and present the framework developed in Doval and Skreta (2022) to capture such complex dynamic games while sidestepping technical mathematical challenges.

In a nutshell, the difficulty arises because mechanisms are part of the histories, and it is impossible to define a measurable structure on them that both allows deriving expected payoffs at given histories and proving a revelation principle. Measurability plays two key roles: it enables evaluating the principal's payoff under a given strategy and comparing it to the payoff from any deviation, which determines sequential rationality. Standard workarounds from the competing-principals literature (Attar et al., 2021)—such as restricting to bounded Borel classes of mechanisms—fail here because such classes are not closed under composition, which is essential for proving revelation principle-style results.[19]

To address these issues, Doval and Skreta (2022) introduce a framework for mechanism design with limited commitment and continuum type spaces that avoids defining the principal's and agent's strategies as measurable functions of past mechanisms. In this new framework, the principal offers *agent extensive forms* $\Gamma(\mu_t,a^{t-1},(\varphi_\tau)_{\tau\geq t})$, prescribing output messages and current allocations in each period as a function of the agent's current input messages and participation decisions. Given an agent extensive form, the prior over Θ together with the agent's strategy results in a distribution over outcomes (joint distributions over types and allocations) and associated expected payoffs. The extensive form is indexed by beliefs, past allocations, and output messages, and at each point we can test the designer's sequential rationality by evaluating her payoff from a deviation to another $\tilde{\Gamma}(\mu_t, a^{t-1}, (\varphi_\tau)_{\tau\geq t})$. For full details see Definition 3 through Definition 5 in Section 4.2 in Doval and Skreta (2022).

4.3 Imperfect persistence

In this subsection, we allow for imperfect persistence in the agent's private information—an important extension for applications in which types evolve over time.[20] As in

[19] Attar et al. (2025b) develop a different approach that imposes no constraints on the principals' mechanisms but restricts the agents' equilibrium strategies to belong to Young classes.

[20] Important applications with imperfectly persistent types include dynamic public finance (Golosov et al., 2006; Stantcheva, 2020; Farhi et al., 2012; Golosov and Iovino, 2021) and sequential screening

the previous two subsections, the material here draws heavily from previous joint work with Laura Doval, in particular Doval and Skreta (2024b).

We overview the results in Doval and Skreta (2024b), who follow the literature on dynamic mechanism design—see Pavan et al. (2014)—and consider a *Markov* environment, defined by two properties. First, the agent's private information is described by a non-homogeneous Markov process: In each period $t \geq 1$, the agent's type θ_t is drawn from a set of types Θ according to a distribution $F_t(\cdot|\theta_{t-1}, a_{t-1})$, where (θ_{t-1}, a_{t-1}) denotes the agent's type and allocation in period $t-1$ (with the convention that when $t = 1$, $F_t(\cdot|\theta_{t-1}, a_{t-1}) \equiv F_1$). Second, the principal's and the agent's payoffs are time separable and their period-t flow payoffs depend only on the current allocation and the agent's period-t type. Formally, letting $(a^T, \theta^T) \in (A \times \Theta)^T$ denote the allocations and the agent's private information through period T, the principal's and agent's payoffs are given by

$$W(a^T, \theta^T) = \sum_{t=1}^{T} \delta^t w_t(a_t, \theta_t), \quad U(a^T, \theta^T) = \sum_{t=1}^{T} \delta^t u_t(a_t, \theta_t).$$

Extending the framework of Doval and Skreta (2022) to environments in which the agent's type evolves over time introduces two primary challenges. The first concerns the specification of the mechanism's input. When the agent's type changes dynamically, her private history at time t is given by $\theta^t = (\theta_1, \dots, \theta_t)$. The result in Doval and Skreta (2022) then implies that the designer must elicit the entire vector θ^t, leading to a message space that grows with time.[21] The Markov assumption implies that the agent's reporting incentives depend solely on her current type θ_t, which permits a simpler version of the revelation principle that is particularly well-suited for applications. The second challenge concerns the mechanism's outputs. Since the agent's type evolves over time, there are two natural candidates for the principal's posterior beliefs: the belief at the end of period t about the type history θ^t and the belief at the beginning of period $t+1$ about the updated type $\theta^{t+1} = (\theta^t, \theta_{t+1})$.

The main theorem in Doval and Skreta (2024b) (reproduced as Theorem 2 below) resolves these issues and establishes an analog of the revelation principle of Doval and Skreta (2022) for Markov environments. First, it establishes that the designer only needs to elicit the agent's *current* type θ_t, rather than her entire history of types θ^t.

(Deb and Said, 2015).

[21] The revelation principle in Doval and Skreta (2022) has the agent report his type in every period, making it conceptually closer to Townsend (1988) than to Myerson (1982, 1986). In the latter, under the assumption of fully persistent types, the agent reports only once at the beginning, and incentive compatibility is guaranteed only along truthful paths. By contrast, repeated reporting in Townsend (1988) ensures incentive compatibility both on and off the equilibrium path.

Second, the theorem identifies the principal's belief about θ_t at the end of period t as the canonical output message of the mechanism. Coupled with the Markov assumption, this implies that it suffices to track the principal's belief over the current type θ_t, rather than over the full type history θ^t. These simplifications are particularly advantageous for applications, as illustrated in Doval and Skreta (2025), which explores an application in industrial organization.

Theorem 1 in Doval and Skreta, 2024b identifies a specific mechanism-selection game and a corresponding class of assessments that together can replicate any equilibrium payoff vertor in $G_{\mathcal{I}}$, for any collection $\mathcal{I}$ of input and output message spaces. We refer to this extensive-form game as the *canonical game* and the associated class of assessments as *canonical assessments*.

Canonical game The canonical game is a mechanism-selection game where the principal is restricted to offering mechanisms with *current* type reports as inputs and beliefs over current types as outputs. Formally, only mechanisms with $M = \Theta$ and $S = \Delta(\Theta)$ are permitted. Let $\mathcal{E}^*$ denote the set of equilibrium payoffs in the canonical game.

Definition 6 (Flow Direct Blackwell Mechanisms Doval and Skreta (2024b)). *A mechanism $(\Theta_t, \Delta(\Theta_t), \varphi)$ is a flow direct Blackwell mechanism (f-DBM) if the map $\varphi : \Theta_t \to \Delta(\Delta(\Theta_t) \times A_t)$ can be written as the composition of two mappings:*

$$\beta : \Theta_t \to \Delta(\Delta(\Theta_t)), \quad \alpha : \Delta(\Theta_t) \to \Delta(A_t),$$

which are, respectively, the mechanism's disclosure and allocation rules.

As is the case with DBMs, in an f-DBM the allocation is drawn *independently* of the agent's report, conditional on the output message. We denote the set of all such mechanisms by $\mathcal{M}_C$.

Canonical assessments A *canonical assessment* specifies behavior for the principal and the agent that is, in a certain sense, straightforward.[22] First, the principal restricts attention to f-DBMs. Second, the agent responds to the principal's equilibrium mechanism by choosing to participate and report truthfully. Third, messages are interpreted *literally*: upon participation, the agent truthfully reports her current type, and any output message $\mu \in \Delta(\Theta)$ reflects the principal's updated belief about the agent's current type. Formally, letting $h^{t+1} = (h^t, \mathbf{M}_t, \mu, a_t)$, the marginal of $\mu_{t+1}(h^{t+1})$ on θ_t coincides with μ, i.e., $\mu_{t+1\Theta_t}(h^{t+1}) = \mu$. Finally, the agent's strategy depends only on her period-t type and the public history. We denote by $\mathcal{E}^C$ the set of equilibrium pay-

[22]See Doval and Skreta (2024b) for full details.

offs under canonical perfect Bayesian equilibrium (PBE) assessments in the canonical game, and by $\mathcal{E}^*_{\mathcal{I}}$ the set of equilibrium payoffs in the mechanism selection game $G_{\mathcal{I}}$.

Theorem 2 (Revelation principle for Markov environments Doval and Skreta, 2024b)**.** *For any PBE assessment of any mechanism-selection game $G_{\mathcal{I}}$, a payoff-equivalent canonical PBE of the canonical game exists. That is,*

$$\bigcup_{\mathcal{I}} \mathcal{E}^*_{\mathcal{I}} = \mathcal{E}^* = \mathcal{E}^C.$$

At first glance, this theorem appears analogous to Theorem 1 in Doval and Skreta, 2022. However, the latter replicates *outcomes*—that is, joint distributions $\Delta(\Theta \times A^T)$—and is therefore stronger than Theorem 2 (which corresponds to Theorem 1 in Doval and Skreta, 2024b), which replicates the principal's and the agent's payoffs. To replicate outcomes, one must elicit the agent's entire type history, whereas flow DBMs are simpler in that they elicit only the agent's *current* type. It is the Markov environment assumption that justifies restricting attention to f-DBMs without loss of generality. The simpler type report comes at the small cost of replicating payoffs only rather than outcomes.

By Theorem 1 in Doval and Skreta, 2024b the analysis of equilibrium payoffs in any mechanism-selection game can, without loss of generality, be reduced to studying canonical equilibria in the canonical game. This simplification streamlines the characterization of equilibrium outcomes and reduces the principal-optimal equilibrium to a constrained optimization problem. Doval and Skreta (2025) apply this to study how firms' data retention practices shape product design and pricing, deriving a dynamic envelope condition in limited-commitment settings that parallels that in Pavan et al. (2014). They show that personalization enables price discrimination since product choice reveals consumer information. To mitigate ratcheting effects, firms may restrict product variety—pooling consumers to limit learning. Thus, the inability to commit not to use past data introduces a novel distortion by reducing available varieties.

5 Putting the revelation principle to work

This section applies the revelation principle of Doval and Skreta (2022) to dynamic principal–agent settings with finite or continuum type spaces. We revisit the classic durable good context to illustrate how the aforementioned result simplifies the analysis.

Direct Blackwell mechanisms allow us to separately design the disclosure rule and the allocation rule. This decomposition cuts through the Gordian knot of belief dependence on past and future mechanisms, as we now proceed to illustrate. We solve the principal's problem in two parts: first, find the optimal allocation rule for a fixed disclosure rule, and then optimize over disclosure rules. The literature on information design teaches us that the problem of finding the optimal disclosure β can be transformed into one of finding the optimal distribution over posteriors. In other words, to optimally design *Direct Blackwell mechanisms*, one combines tools from the mechanism design and information design literatures and solves a constrained information design problem. Doval and Skreta (2024a) provide general tools to analyze information design problems subject to such constraints, which we overview in Section 5.3.

5.1 Sale of a durable good: binary types; $T = 2$

This section draws on presentations and notes from joint work with Laura Doval. We consider a two-period interaction between a seller and a buyer over the sale of a single unit of a durable good. The seller assigns zero value to the good and faces a buyer whose private type is denoted by $\theta \in \Theta \equiv \{\theta_L, \theta_H\}$, with prior belief $\mu_1 = \Pr(\theta = \theta_H)$. An allocation in each period is given by a pair $(q, x) \in \{0,1\} \times \mathbb{R}$, where q indicates whether the good is sold ($q = 1$) or not ($q = 0$), and x is a monetary transfer from the buyer to the seller. If the good is sold in the first period, the game terminates. Given a vector of period-1 and period-1 allocations $\{(q_t, x_t)\}_{t=1}^{2}$, the buyer's payoff is

$$U(\cdot, \theta) = \sum_{t=1}^{2} \delta^{t-1}(q_t\theta - x_t),$$

and the seller's payoff is

$$W(\cdot, \theta) = \sum_{t=1}^{2} \delta^{t-1} x_t,$$

where $\delta \in (0,1)$ is a common discount factor.

Timing Structure At the beginning of each period $t \in \{1,2\}$, the seller offers a mechanism to the buyer. The buyer then decides whether to accept or reject the offer. If the buyer accepts, he proceeds to privately report in the mechanism, and an allocation—consisting of a transfer and possibly the delivery of the good—is realized based on the public output message triggered by the buyer's report. The output message and the allocation are publicly observed. If the good is traded, the game ends immediately. Otherwise, the game proceeds to period $t + 1$. If the buyer rejects the mechanism, the default allocation $(q, x) = (0, 0)$ is implemented, and the game similarly moves to the next period.

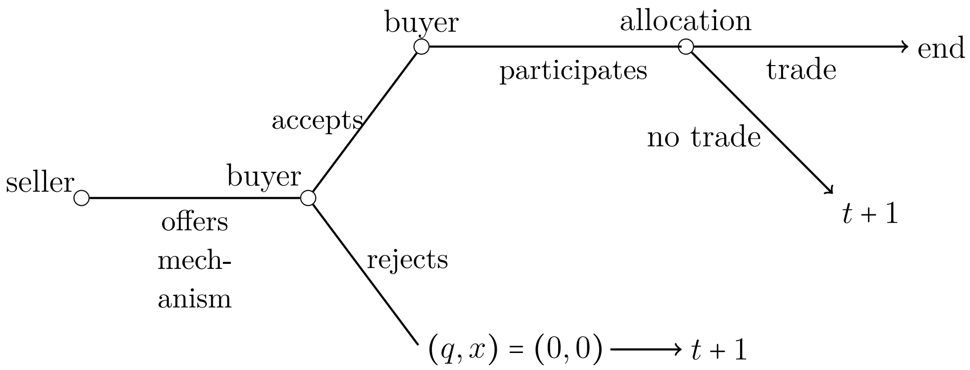

Figure 6: Timing in each period of the durable good sale game

Period 2: Static Optimal Mechanism In the final period of the game, the seller has full commitment power. Consequently, the standard revelation principle applies, and the mechanism design problem reduces to a static screening problem. Let μ_2 denote the seller's posterior belief that the buyer is of high type, i.e., $\mu_2 = \Pr(\theta = \theta_H)$.

The optimal mechanism is characterized by a cutoff belief $\bar{\mu} \equiv \theta_L/\theta_H$. When $\mu_2 < \bar{\mu}$, the seller optimally sells the good to both types at a price equal to θ_L. When $\mu_2 > \bar{\mu}$, it is optimal to sell only to the high type at a price equal to θ_H, thereby extracting full surplus from that type while excluding the low type. This behavior is illustrated in Figure 7, where the mechanism transitions from pooling to separation at $\mu_2 = \bar{\mu}$.

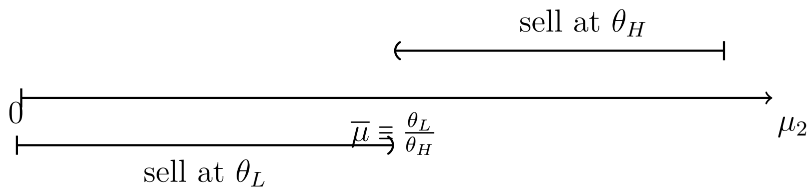

Figure 7: Allocation regions as a function of posterior belief μ_2

At $\mu_2 = \overline{\mu}$, both prices are optimal. (This multiplicity will be leveraged below when we consider what is optimal from the perspective of period-1 seller, see Figure 9 below). Cutoff $\overline{\mu}$ reflects the rent left to the high type under pooling. Whenever the seller sells to both types at price θ_L, the high type receives informational rents. These rents become increasingly costly as μ_2 grows. At the threshold $\mu_2 = \bar{\mu}$, the seller is indifferent: the marginal rent to the high type exactly offsets the marginal gain from serving the low type. Formally, at $\mu_2 = \bar{\mu}$ we have $\hat{\theta}_L(\mu_2) = 0$, where $\hat{\theta}_L(\mu_2) \equiv \theta_L - \frac{\mu_2}{1-\mu_2}(\theta_H - \theta_L)$ denotes the virtual valuation of the low type under belief μ_2: Letting $\Delta\theta = \theta_H - \theta_L$, we can express θ_L thought the following identity:

$$\theta_L = \mu_2(\theta_H - \Delta\theta) + (1-\mu_2)\theta_L = \mu_2\theta_H + (1-\mu_2)\left(\theta_L - \frac{\mu_2}{1-\mu_2}\Delta\theta\right) = \mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_2).$$

The last equality makes explicit how the virtual type $\hat{\theta}_L(\mu_2)$ governs the trade-off. Let $R_2(\mu_2)$ denote the seller's optimal expected revenue in period 2. From the above analysis it follows that it can be expressed as:

$$R_2(\mu_2) = \begin{cases} \theta_L = \mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_2), & \text{if } \mu_2 \le \bar{\mu}, \\ \mu_2\theta_H, & \text{if } \mu_2 > \bar{\mu}. \end{cases}$$

This revenue function is piecewise smooth and is illustrated in Figure 8. The function is flat when both types are served (the period-2 price is θ_L) and becomes strictly increasing once only the high type is targeted (the period-2 price is θ_H).

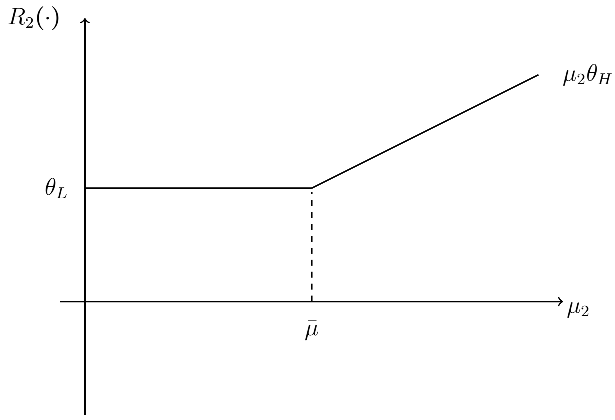

Figure 8: Seller's revenue in period 2 as a function of belief μ_2

Period 1: Direct Blackwell mechanisms A Direct Blackwell mechanism consists of a communication device β mapping a type report to a distribution over posteriors and an allocation rule mapping each posterior to a distribution over trade decisions and transfers. In period 1, an important simplification arises from the combination of quasilinear utility and the separation between allocation and information. Specifically, there is no need to randomize monetary transfers. This simplification holds generally in quasilinear environments (see Doval and Skreta, 2024a and Section 5.3 below). Given any output message μ_2, the transfer function $x(\mu_2)$ denotes the expected payment the buyer makes, and $q(\mu_2)$ denotes the probability that the good is sold.

The seller's optimal outcome is characterized by a constrained mechanism design problem, which we refer to as Program (OPT), described as follows:

$$R_1(\mu_1) \equiv \max_{\beta, q, x} \sum_{\mu_2 \in \Delta(\Theta)} \left(\sum_{\theta \in \Theta} \mu_1(\theta)\beta(\mu_2 \mid \theta) \right) [x(\mu_2) + (1 - q(\mu_2))\delta R_2(\mu_2)], \qquad \text{(OPT)}$$

subject to the following participation, truth-telling and Bayes' plausibility constraints:

$$\sum_{\mu_2 \in \Delta(\Theta)} \beta(\mu_2 \mid \theta) \left[\theta q(\mu_2) - x(\mu_2) + (1 - q(\mu_2))\delta u^*(\mu_2, \theta)\right] \geq 0 \quad \forall \theta \in \{\theta_L, \theta_H\} \tag{PC}$$

$$\sum_{\mu_2 \in \Delta(\Theta)} \left(\beta(\mu_2 \mid \theta) - \beta(\mu_2 \mid \theta')\right) \left[\theta q(\mu_2) - x(\mu_2) + (1 - q(\mu_2))\delta u^*(\mu_2, \theta)\right] \geq 0 \quad \forall \theta, \theta' \in \Theta \tag{TT}$$

$$\mu_2(\theta_H) \sum_{\theta \in \Theta} \mu_1(\theta)\beta(\mu_2 \mid \theta) = \mu_1(\theta_H)\beta(\mu_2 \mid \theta_H) \quad \forall \mu_2 \in \Delta(\Theta), \tag{BP}$$

where $u^*(\mu_2, \theta)$ is the buyer's continuation payoff when his type is θ and the seller's posterior in μ_2. These constraints ensure that the buyer is willing to participate (PC), report truthfully (TT), and that the induced posteriors are consistent with Bayes' rule (BP).

Dynamic virtual surplus representation At the optimal mechanism, standard arguments imply that the participation constraint binds for the low-valuation buyer, meaning that the seller extracts all surplus from type θ_L. Additionally, the truth-telling constraint binds for the high-valuation buyer, who is indifferent between reporting truthfully and mimicking the low type. Leveraging these binding constraints, the seller's problem can be reformulated as a virtual surplus maximization problem. Let $\tau(\mu_2)$ denote the total probability that the posterior μ_2 is induced (i.e., the marginal of the communication device β), and let $q(\mu_2)$ denote the probability of trade conditional on μ_2. The seller's expected revenue can then be expressed as the following *relaxed* program:

$$R_1(\mu_1) = \max_{\tau, q} \sum_{\mu_2 \in \Delta(\Theta)} \tau(\mu_2) \left[q(\mu_2)\left(\mu_2 \theta_H + (1 - \mu_2)\hat{\theta}_L(\mu_1)\right) + (1 - q(\mu_2))\delta R(\mu_2; \mu_1)\right],$$

subject to a Bayes' plausibility constraint

$$\sum_{\mu_2 \in \Delta(\Theta)} \tau(\mu_2)\mu_2(\theta_H) = \mu_1(\theta_H).$$

Here, $\hat{\theta}_L(\mu_1) = \left(\theta_L - \frac{\mu_1}{1-\mu_1}\Delta\theta\right)$ denotes the virtual type of the low-valuation buyer as perceived by the seller under belief μ_1. The term $R(\mu_2; \mu_1)$ denotes the seller's expected continuation payoff in period 2 following posterior μ_2 evaluated from the period 1 perspective at which point the seller pays rents to the high type with probability μ_1 rather than μ_2. The term inside the square brackets captures the expected virtual surplus: when trade occurs ($q(\mu_2) = 1$), the seller receives a virtual valuation based on the mix of buyer types; when trade is delayed ($q(\mu_2) = 0$), the seller discounts the expected future revenue. The program is relaxed because it only reflects local incentive compatibility constraints rather than global incentive constraints reflected in *monotonicity constraints*, see (M) in Section 5.3.

The relaxed program highlights how dynamic incentive constraints distort the seller's problem toward selecting information structures that maximize ex-ante expected revenue, subject to Bayes' plausibility. The continuation value from period 2, evaluated in period 1, depends on the posterior belief μ_2 and the prior μ_1 see Figure 9. Specifically, for a given posterior μ_2, the seller's discounted expected revenue, henceforth *value of delay* is

$$\delta R_2(\mu_2;\mu_1) = \begin{cases} \delta\left(\mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_1)\right), & \text{if } \mu_2 < \bar{\mu}, \\ \delta\mu_2\theta_H, & \text{if } \mu_2 > \bar{\mu}. \end{cases}$$

The shape of the value of delay $\delta R_2(\mu_2;\mu_1)$ varies depending on whether the prior μ_1 lies below or above the cutoff $\bar{\mu}$. This is illustrated in the two-panel figure below.

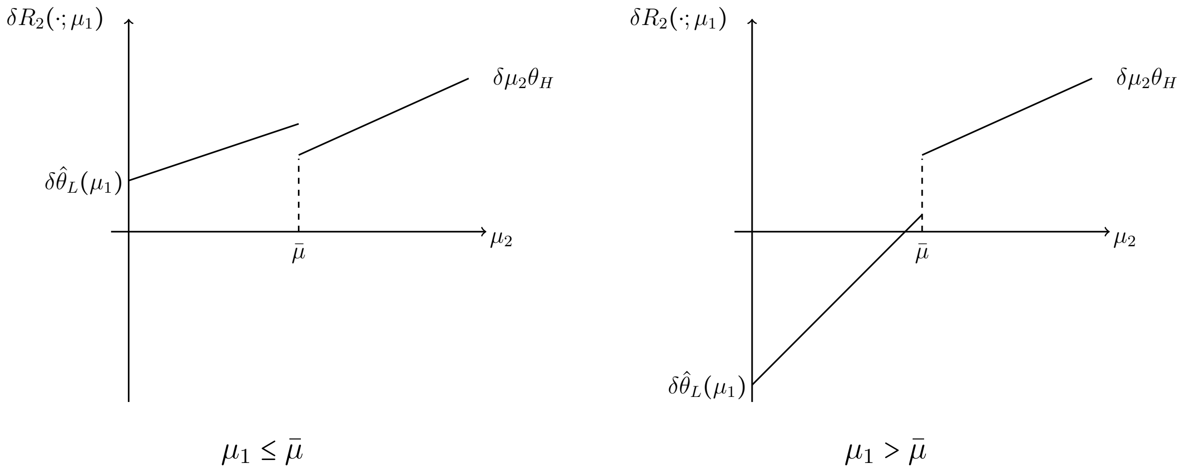

Figure 9: Intertemporal conflict: seller's revenue in period 2 as a function of belief μ_2 from the perspective of period 1 (and belief μ_1)

Constrained information design in period 1 The seller's optimization problem in period 1 can be represented as a dynamic virtual surplus maximization problem over distributions of posteriors and allocation rules. Formally, the seller selects $\tau \in \Delta\Delta\Theta$ and $q : \Delta(\Theta) \to [0,1]$ to solve:

$$\max_{\tau,q} \sum_{\mu_2 \in \Delta(\Theta)} \tau(\mu_2)\left[q(\mu_2)\left(\mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_1)\right) + (1-q(\mu_2))\delta R_2(\mu_2;\mu_1)\right],$$

where $\tau(\mu_2)$ is the probability of reaching posterior μ_2, and $q(\mu_2)$ is the probability of trade at that posterior (the transfers are already replaced out using the binding constraints mentioned above).

The decision whether to trade or delay hinges on comparing two quantities: the virtual surplus from trade, and the value of delay. At each posterior μ_2, the seller optimally sets $q(\mu_2) = 1$ if the virtual surplus from trade exceeds the value of delay, and $q(\mu_2) = 0$ otherwise. This allows the objective to be rewritten as:

$$\max_{\tau \in \Delta\Delta\Theta} \sum_{\mu_2 \in \Delta(\Theta)} \tau(\mu_2) \max\left\{\mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_1),\, \delta R_2(\mu_2;\mu_1)\right\}.$$

This structure reflects a classic information design problem: the seller chooses a distribution over posteriors τ to optimize a pointwise maximum of two value functions. A solution concavifies the objective and for the scenario depicted in Figure 10 where $\mu_1 > \overline{\mu}$ places mass on at most two posteriors, typically $\mu_2 = \bar{\mu}$ and $\mu_2 = 1$. The seller sells the good when $\mu_2 = 1$ and delays trade when $\mu_2 = \bar{\mu}$. In both cases, the posted price is θ_H, and the mechanism resembles a dynamic version of posted pricing.

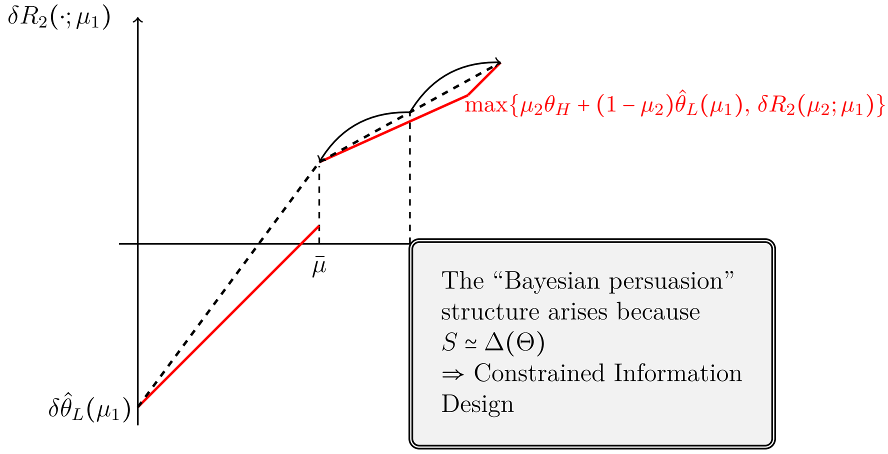

Figure 10: Optimal disclosure and trade decision in period 1

This analysis reveals a core trade-off in dynamic mechanism design with limited commitment: the seller must choose between tailoring the allocation to the agent's current report and designing information structures that elicit valuable signals about the agent's type. Unlike in static settings with commitment, acquired information cannot be discarded ex post. Thus, the designer must manage the long-term consequences of present disclosures. In dynamic contexts, the seller coordinates across periods via these posteriors, which play a dual role as public messages to determine current allocations and as determinants of future mechanisms.

The constraints in OPT yield a tractable representation of the buyer's behavior and its informational implications for $t = 2$. This allows us to sidestep the complexity of mixed strategies (Laffont and Tirole, 1988; Bester and Strausz, 2001) and instead recast the problem of identifying the seller-optimal PBE outcome as solving the program OPT, which merges insights from information and mechanism design. In this example, the solution of the relaxed problem did indeed satisfy monotonicity, and given that it involved two posteriors, we could, as usual in mechanism design, construct the transfers from the allocation rule. However, it is possible that even though the solution of the relaxed program satisfies (M), there may not be enough posteriors to construct transfers; see Doval and Skreta (2024a) for such an example.

When the type space Θ is finite, Doval and Skreta (2022) (see Proposition 1) show that the program (OPT) admits a further simplification: the seller can, without loss of generality, restrict attention to mechanisms that induce finitely supported distributions over posteriors. This result relies on the fact that the disclosure rule β satisfies a Bayes plausibility condition, which implies that the distribution over posteriors lies in a convex set. By Carathéodory's theorem (Rockafellar, 1970), the seller can then achieve the same outcomes using only a finite set of posteriors. In particular, the number of posteriors need not exceed the number of states and of binding constraints. See also Proposition 3 in Bester and Strausz (2007). Doval and Skreta (2024a) establish that further simplifications are possible in some fairly general finite-type settings.

The analysis in this section revealed that at the binary-state example we studied the number of posteriors at the optimum was just two: one belief corresponding to $q = 1$, (sell belief) and one belief corresponding to $q = 0$ (delay belief). One might be tempted to conclude that the result is true regardless of the number of types. However, it is not. We discuss how to derive the analogue of OPT for the case of continuum of types in Section 5.2.

5.2 Sale of a durable good: continuum of types; $T = 2$.

We revisit the example of Section 5.1 under the assumption that types are drawn from a continuum, $\Theta = [\underline{\theta}, \overline{\theta}] \subseteq \mathbb{R}$, with $\underline{\theta} < \overline{\theta}$. We use F_1 and F_2 to denote the seller's prior and posterior beliefs, respectively. This illustration will show how one can use the framework to formalize an analogue of OPT for a continuum of types and to derive a dynamic virtual representation of the designer's expected payoff.

Seller's program. By the analogue of Theorem 1 for a continuum of types (see Theorem 2 in Doval and Skreta, 2022), it is without loss to restrict the seller to DBMs and analyze the associated canonical agent-PBE. Program OPT continues to represent the seller's optimal revenue over PBE-feasible outcomes. The formulation of the mechanism selection game for a continuum of types (see Definition 3–Definition 5 in Doval and Skreta, 2022) requires that the period-2 outcome distribution be PBE-feasible given the seller's posterior belief F_2, which in this setting implies that the seller's period-2 mechanism maximizes revenue given F_2. By Proposition 2 in Skreta (2006), this is a posted price $p_2(F_2)$ (note that F_2 can have atoms and gaps in the support).

Dynamic virtual surplus representation With a continuum of types, we can apply the envelope theorem to eliminate transfers and express the seller's objective in terms of virtual surplus, reducing the seller's program (the analogue of OPT for a continuum of types) to:

$$\max_{P_{\Delta(\Theta)},q} \int_{\Delta(\Theta)} \Bigg[q(F_2) \int_{\underline{\theta}}^{\overline{\theta}} \Big(\theta - \frac{1-F_1(\theta)}{f_1(\theta)}\Big) F_2(d\theta) \\ + (1-q(F_2))\delta \int_{p_2(F_2)}^{\overline{\theta}} \Big(\theta - \frac{1-F_1(\theta)}{f_1(\theta)}\Big) F_2(d\theta) \Bigg] P_{\Delta(\Theta)}(dF_2) \tag{2}$$

subject to (i) $P_{\Delta(\Theta)}$ being Bayes' plausible given F_1, and (ii) a monotonicity condition requiring that higher types trade with weakly higher probability in expectation.

In this formulation, if the seller trades at posterior F_2, he earns the expected virtual surplus (under F_2, but computed with F_1's virtual values). If not, he receives the discounted virtual surplus from posting price $p_2(F_2)$. Since $p_2(F_2)$ is optimal relative to F_2, but the seller values outcomes using F_1, this can create misalignment between the seller's period-1 and period-2 objectives as we saw already for the binary-type setting.[23]

Equation 2 defines an information design problem with a continuum of states in which the designer's payoff depends on the entire shape of the distribution (not just the posterior mean or other moments). Consequently, it falls outside the scope of existing methods in information design (cf. Kolotilin, 2018; Dworczak and Martini, 2019). Nevertheless, the virtual-surplus representation of the seller's problem sheds new light to the sale of a durable good under limited commitment. Skreta (2006) shows that, among the transparent indirect mechanisms (which include mechanisms in Bester and Strausz, 2001 but also allow for randomizations), posted prices are optimal for the seller. In contrast, when the seller has access to direct Blackwell mechanisms, Proposition 2 in Doval and Skreta (2022) establishes conditions under which the seller profitably deviates from the optimal posted-price mechanism to a mechanism that induces three beliefs: one that corresponds to trade ($q = 1$), one to delay ($q = 0$), and one to rationing ($0 < q < 1$). Hence, the outcome distribution induced by posted prices

[23]If, for example, the posterior is such that the posterior virtual value $\theta - \frac{1-F_2(\theta)}{f_2(\theta)}$ is well-defined and increasing, then $p_2(F_2)$ sets the posterior rather than the prior virtual value to zero, that is, $p_2(F_2) - \frac{1-F_2(p_2(F_2))}{f_2(p_2(F_2))} = 0$ instead of $p_2(F_2) - \frac{1-F_1(p_2(F_2))}{f_1(p_2(F_2))} = 0$.

is not PBE-feasible. The solution to the program in Equation 2 remains elusive.

5.3 Limited commitment as constrained information design

This section draws heavily from Doval and Skreta (2024a) who study abstract constrained information design problems to discuss how Theorem 3.1 in that paper can be leveraged to inform the characterization of optimal mechanisms in settings in which the principal has limited commitment. Expressing the principal's problem as a constrained information-design program (see OPT_{LC}) highlights the conceptual link between limited-commitment mechanism design and information design. The principal effectively acts as a sender who designs an information structure for his future self. Unlike in standard persuasion models (Kamenica and Gentzkow, 2011; Taneva, 2019), however, she must also determine the distribution of period-1 allocations associated with each posterior, and the set of feasible posteriors is constrained by the agent's incentive and participation requirements, which themselves depend on the induced future allocations.

Consider a principal who interacts with a privately informed agent over two periods, $t \in \{1, 2\}$. The timing of the game and the space of mechanisms for the principal are as described in Section 4, with the difference that the allocation $a_2 \in A_2$ lends itself to multiple interpretations. For instance, following Bester and Strausz (2001) and Bester and Strausz (2007), it can capture in reduced form the principal's limited commitment. Alternatively, it can capture a continuation equilibrium's expected discounted allocation in a longer-horizon interaction as in Skreta (2006) (but with the difference that a_2 does not depend on θ; whereas expected discounted allocations can, in general, vary with θ).

As a reminder, we let $\mu_1 \in \Delta(\Theta)$ denote the principal's prior belief about the agent's type θ. In each period t, as a result of the interaction, an allocation $a_t \in A_t$ is determined, where A_t is the set of allocations in period t. There is a correspondence $\mathscr{A} : A_1 \rightrightarrows A_2$ that describes the set of feasible period 2 allocations as a function of the allocation in period 1. Let $W(a^T, \theta) = W(a_1, a_2, \theta)$ and $U(a_1, a_2, \theta)$ denote the principal and the agent's payoff, respectively, when the allocation is (a_1, a_2) and the type is θ. As in Section 2.1, we assume an allocation (a_1^*, a_2^*) exists such that $U(a_1^*, a_2^*, \theta) = 0$ for all $\theta \in \Theta$. This allocation plays the role of the outside option in what follows.

At the principal-optimal Perfect Bayesian equilibrium, the period-2 allocation must be sequentially rational. For each $a_1 \in A_1$, let $\mathcal{O}_2^*(a_1, \mu)$ denote the set of feasible continuation equilibrium allocations in period 2 when the principal's belief about θ is

μ. From the perspective of the period-1 principal, the optimal continuation is then

$$a_2^*(a_1,\mu) \equiv \arg\max_{a_2 \in \mathcal{O}_2^*(a_1,\mu)} \sum_{\theta\in\Theta} \mu(\theta) W(a_1,a_2,\theta). \tag{3}$$

The assumptions on $\mathcal{O}_2^*(a_1,\mu)$ in footnote 12 in Doval and Skreta (2024a) ensure that this problem is well-defined.

Leveraging Theorem 1 in Doval and Skreta (2022) we can write the principal's problem as follows:

$$\max_{\beta:\Theta\mapsto\Delta\Delta(\Theta),\alpha:\Delta(\Theta)\mapsto\Delta(A_1),a_2\in\mathcal{O}_2^*} \sum_{\theta\in\Theta} \mu_1(\theta)\mathbb{E}_{\beta(\cdot|\theta)}\left[\mathbb{E}_{\alpha(\cdot|\mu)}\left[W(a_1,a_2(a_1,\mu),\theta)\right]\right] \tag{$\mathscr{P}$}$$

$$\text{s.t.} \begin{cases} (\forall\theta\in\Theta) & \mathbb{E}_{\beta(\cdot|\theta)}\left[\mathbb{E}_{\alpha(\cdot|\mu)}\left[U(a_1,a_2(a_1,\mu),\theta)\right]\right] \geq 0 \\ (\forall\theta\in\Theta)(\forall\overline{\theta}\neq\theta) & \mathbb{E}_{\beta(\cdot|\theta)-\beta(\cdot|\overline{\theta})}\left[\mathbb{E}_{\alpha(\cdot|\mu)}\left[U(a_1,a_2(a_1,\mu),\theta)\right]\right] \geq 0 \end{cases},$$

where the two sets of constraints are the agent's participation and truthtelling constraints. Furthermore, the disclosure rule β must satisfy that for all measurable subsets $\widetilde{U}$ of $\Delta(\Theta)$ and all subsets $\widetilde{\Theta}$ of Θ,

$$\sum_{\overline{\theta}\in\widetilde{\Theta}} \beta(\widetilde{U}|\overline{\theta})\mu_1(\overline{\theta}) = \sum_{\theta\in\Theta} \int \mu(\widetilde{\Theta})\beta(d\mu|\theta)\mu_1(\theta). \tag{4}$$

When for each θ, the distribution $\beta(\cdot|\theta)$ has finite support, (4) simplifies to:

$$\mu(\theta) = \frac{\mu_1(\theta)\beta(\mu|\theta)}{\sum_{\overline{\theta}\in\Theta}\mu_1(\overline{\theta})\beta(\mu|\overline{\theta})}. \tag{5}$$

As explained above, Doval and Skreta (2022, Proposition 1) show that with finitely many types, finite support is without loss of generality. The results that follow provide tighter bounds on the number of required posteriors.

From disclosure rules to distributions over posteriors As established in the information design literature, the problem of finding the optimal disclosure β is equivalent to finding the optimal distribution over posteriors. Any disclosure rule β induces a distribution $\tau \in \Delta\Delta(\Theta)$ over posteriors such that $\beta(\widetilde{U}|\theta)\mu_1(\theta) = \int_{\widetilde{U}} \mu(\theta)\tau(d\mu)$ for all $\theta\in\Theta$ and measurable subsets $\widetilde{U} \subset \Delta(\Theta)$. See Doval and Skreta (2024a) for full details. We can then recast the principal's objective and the agent's constraints as expectations over Bayes' plausible distribution over posteriors (that is, $\tau \in \Delta_{\mu_1}\Delta(\Theta)$), and reformulate program $\mathscr{P}$ as a constrained information design problem:

$$\max_{\tau \in \Delta_{\mu_1}\Delta(\Theta), \alpha:\Delta(\Theta)\mapsto\Delta(A_1), a_2 \in \mathcal{O}_2^*} \mathbb{E}_\tau \left[\mathbb{E}_{\alpha(\cdot|\mu)} \left[\sum_{\theta\in\Theta} \mu(\theta) W(a_1, a_2(a_1,\mu), \theta) \right] \right] \qquad (\text{OPT}_{LC})$$

$$s.t. \begin{cases} (\forall \theta \in \Theta) & \mathbb{E}_{\tau(\cdot)} \left[\mathbb{E}_{\alpha(\cdot|\mu)} \left[\frac{\mu(\theta)}{\mu_1(\theta)} U(a_1, a_2(a_1,\mu), \theta) \right] \right] \geq 0 \\ (\forall \theta \in \Theta)(\forall \overline{\theta} \neq \theta) & \mathbb{E}_{\tau(\cdot)} \left[\mathbb{E}_{\alpha(\cdot|\mu)} \left[\left(\frac{\mu(\theta)}{\mu_1(\theta)} - \frac{\mu(\overline{\theta})}{\mu_1(\overline{\theta})} \right) U(a_1, a_2(a_1,\mu), \theta) \right] \right] \geq 0 \end{cases}$$

which corresponds to a special case of program CID of Doval and Skreta (2024a).

We next focus on the case where the agent's preferences exhibit *increasing differences*, adapted to account for lotteries over allocations:

Definition 7 (Bester and Strausz, 2007; Celik, 2015; Kartik et al., 2024)**.** *The family* $\{U(\cdot,\theta) : \theta \in \Theta\}$ *satisfies* monotonic expectational differences *if for any two distributions* $P, Q \in \Delta(A_1 \times A_2)$*,* $\int U(\cdot,\theta_i) d(P-Q)$ *is monotone in* i *which holds if, and only if* U *takes the form*

$$U(a_1, a_2, \theta_i) = b(\theta_i) h_1(a_1, a_2) + h_2(a_1, a_2) + c(\theta_i),$$

where h_1, h_2 *are finitely integrable and* b *is monotonic.*

Without loss of generality, assume that b is weakly increasing, so that θ_1 is the agent's "lowest type." Moreover, assuming that $h_1(a_1^*, a_2^*) = \min_{(a_1,a_2):a_2\in\Gamma(a_1)} h_1(a_1,a_2)$ ensures that if the lowest type θ_1 participates, all types do. As under commitment, monotonic expectational differences ensure that $\mathscr{P}$ coincides with a simpler program imposing only a subset of the incentive constraints: (i) that the agent's participation constraint holds at θ_1, and (ii) adjacent incentive compatibility constraints are satisfied (Proposition 4.1 in Doval and Skreta, 2024a). That result, also implies that any solution to $\mathscr{P}$ utilizes at most $3N-1$ posteriors (Corollary 4.1 in Doval and Skreta, 2024a) where N is the cardinality of the type space.

Doval and Skreta (2024a) relate their Proposition 4.1 to earlier results showing that adjacent incentive constraints can imply global incentive compatibility. While Bester and Strausz (2007) establish this under specific utility forms, they do not account for participation constraints. Celik (2015) assume our condition in Definition 7 and assert the same implication without proof. Proposition 4.1 provides a complete proof and unifies these results leveraging Kartik et al. (2024), who characterize the utility functions satisfying that condition.

Transferable utility Transferable utility further simplifies the characterization of an optimal mechanism by reducing even further the number of posteriors that the mechanism employs. In such settings, the set of period 1 allocations is $A_1 = A_1' \times \mathbb{R}_+$, where the second coordinate is a payment from the agent to the principal. An element of A_1 is then $a_1 = (a_1', x)$ and we assume that $\mathscr{A}((a_1', x)) = \mathscr{A}(a_1')$. Finally, suppose the agent's and principal's payoffs are given by:

$$W(a_1, a_2, \theta) = \tilde{v}(a_1', a_2, \theta) + x,\ U(a_1, a_2, \theta) = \tilde{u}(a_1', a_2, \theta) - x.$$

As we saw in durable good two-type example in Section 5.1, under transferable utility, restricting attention to mechanisms that do not randomize over transfers entails no loss of generality. Henceforth, we substitute x with its expected value under the mechanism conditional on posterior μ, denoted by $x(\mu)$. We also define $\tilde{\alpha} : \Delta(\Theta) \to \Delta(A_1')$ as the marginal of α on A_1'.

Under monotonic expectational differences and transferable utility, the lowest type's participation constraint binds. Moreover, in this environment, one can obtain a simpler problem if one can establish (as we did in Section 5.1) that, for the particular problem at hand, only downward-looking incentive constraints bind at the optimum. Then, the problem reduces to the corresponding relaxed program:

$$\max_{\tau \in \Delta_{\mu_1}\Delta(\Theta), \tilde{\alpha}:\Delta(\Theta)\mapsto\Delta(A_1'), x:\Delta(\Theta)\mapsto\mathbb{R}, a_2(\cdot)\in\mathcal{O}_2^*(\cdot)} \mathbb{E}_\tau\left[\mathbb{E}_{\tilde{\alpha}(\cdot|\mu)}\left[\sum_{\theta\in\Theta}\mu(\theta)\tilde{v}(a_1', a_2(a_1', \mu), \theta) + x(\mu)\right]\right] \tag{$\mathscr{R}$}$$

$$\text{s.t.}\begin{cases} \mathbb{E}_\tau\left[\mathbb{E}_{\tilde{\alpha}(\cdot|\mu)}\left[\frac{\mu(\theta_1)}{\mu_1(\theta_1)}\left(\tilde{u}(a_1', a_2(a_1', \mu), \theta_1) - x(\mu)\right)\right]\right] = 0 \\ (\forall i \in \{2, \ldots, N\}) \quad \mathbb{E}_\tau\left[\mathbb{E}_{\tilde{\alpha}(\cdot|\mu)}\left[\left(\frac{\mu(\theta_i)}{\mu_1(\theta_i)} - \frac{\mu(\theta_{i-1})}{\mu_1(\theta_{i-1})}\right)\left(\tilde{u}(a_1', a_2(a_1', \mu), \theta_i) - x(\mu)\right)\right]\right] = 0 \end{cases},$$

which is obtained by dropping the *monotonicity constraints*:[24]

$$\mathbb{E}_\tau\left[\mathbb{E}_{\tilde{\alpha}(\cdot|\mu)}\left[\left(\frac{\mu(\theta_i)}{\mu_1(\theta_i)} - \frac{\mu(\theta_{i-1})}{\mu_1(\theta_{i-1})}\right)\left(\tilde{u}(a_1', a_2(a_1', \mu), \theta_i) - \tilde{u}(a_1', a_2(a_1', \mu), \theta_{i-1})\right)\right]\right] \geq 0, \tag{M}$$

for each $i \in \{2, \ldots, N\}$. We can use the binding constraints to substitute the transfers out of the principal's program and obtain that the solution to the relaxed program uses at most N posteriors.

Doval and Skreta (2024a) characterize when the solution to the relaxed program can be implemented as a solution to the principal's full problem. Whereas in standard mechanism design with full commitment any relaxed-program solution satisfying the monotonicity constraints can be implemented via suitable transfers, this is not neces-

[24] The constraints in Equation M are obtained by combining the restriction that θ_i does not want to report θ_{i-1} and θ_{i-1} does not want to report θ_i. Under Definition 7, the binding downward-looking incentive constraints together with the monotonicity constraints imply the local constraints in Proposition 4.1 in Doval and Skreta (2024a).

sarily true under limited commitment. The program is relaxed for two reasons: only adjacent constraints enter and only *average* transfers matter. In particular, because the relaxed program imposes N constraints on transfers but may involve fewer than N posteriors, actual transfers that satisfy all constraints need not exist (see the online appendix of Doval and Skreta, 2024a for such an example). To address this, Doval and Skreta (2024a) establish necessary and sufficient conditions for implementability (see Proposition 6.1 in Doval and Skreta, 2024a).

Moreover, when the monotonicity constraints bind, adding one posterior for each binding constraint suffices for implementation; the resulting solution uses at most $N+B$ posteriors, where B denotes the number of binding monotonicity constraints. These results justify focusing on the relaxed program as a benchmark, while highlighting the additional linear-independence requirement that arises under limited commitment.

Remark 2 (Mechanism design with downstream interactions)**.** *Whereas above we emphasized the limited commitment interpretation, by interpreting $\mathcal{O}_2^*(a_1, \mu)$ in Equation 3 as the continuation equilibrium feasible allocations of a setting where an upstream mechanism designer chooses a mechanism anticipating that a downstream third party observes the mechanism's allocation and takes an action $a_2 \in A_2$ in response. The downstream third party may represent another principal as in Calzolari and Pavan (2006b); Pavan and Calzolari (2009), or an aftermarket as in Calzolari and Pavan (2006a); Dworczak (2020). However, note that Equation 3 specifies a selection that is optimal for the period-1 principal. If there is a unique equilibrium feasible allocation in period 2 (as is in the binary-type durable good problem for all $\mu \neq \bar{\mu} = \frac{\theta_L}{\theta_H}$) then, trivially, a solution of Equation 3 selects the continuation allocation $a_2(a_1, \mu)$ maximizes the expected payoff of the period -2 principal. Otherwise, the PBE outcome that is optimal for the period-1 principal specifies equilibrium play may be suboptimal from the perspective of the downstream principal.*

6 Concluding remarks

The assumption that designers can commit to long-term contracts is often unrealistic. Firms repeatedly interact with customers and adjust terms as they learn; governments revise policies ex post during crises; similar challenges arise in redistribution and social insurance programs. This chapter surveyed mechanism design under limited commitment, emphasizing recent tools for dynamic relationships governed by limited commitment. The revelation principle in Doval and Skreta (2022) offers a novel analytical tool for these environments.

Relaxing the commitment assumption allows researchers to revisit important issues with more realistic models. Classic examples include compensation schemes, debt and mortgage contracts, monetary policy, fiscal policy, and social insurance. Problems of limited commitment are ubiquitous in political science, precisely because current incumbents cannot sign contracts that bind their successors. A methodology that speaks to the design of self-enforcing institutions could thus advance the literature on mechanism design in political economy settings.

Beyond these classical settings, one can rely on DBMs to study among other topics:

1. Vertical contracting: In this context, the period-2 principal—the downstream principal—updates beliefs based on data revealed by the upstream interaction. The downstream principal's decisions affect the upstream principal's payoff, creating incentives to influence the downstream principal's behavior through information revelation. Doval and Skreta (2025) show how an upstream firm coarsens its product line to reduce the informativeness of period-1 choices, thereby disciplining the downstream firm's behavior.

2. Privacy policy design: The period 1 principal can be a data broker as illustrated in Doval and Skreta (2025).

3. Platform design: Platforms continually reoptimize as they learn from repeated consumer interactions. As we enter a new era of AI-assisted decision making, consumers will increasingly rely on sophisticated AI advisors who anticipate how current behavior affects future prices and choices. In other words, consumers may become more aware of how their actions shape future terms—the friction at the heart of this chapter. Hermann and Puntoni (2024) investigate changes on the firm side, but the consumer landscape is rapidly evolving.

Despite this progress, applying the results overviewed in this chapter exposes how little we know about how optimal mechanisms change as we vary the principal's information (comparative statics of the shape of mechanisms with respect to the prior). Beyond the need for mechanism design comparative statics and characterizations of optimal mechanisms for arbitrary priors, we also need new information design tools for continuum type spaces. These tools will further accelerate progress.

Countering the ratchet effect One can also explore tools and institutions that alleviate frictions from the principal's reoptimization under limited commitment.[25]

[25]See also Acharya and Ortner (2017) for a related discussion.

Bisin and Rampini (2006) show that *markets* can mitigate the distortions caused by limited commitment. Specifically, when agents can engage in hidden trades, such trades can act as a commitment device for the principal. In contrast to earlier work (Golosov and Tsyvinski, 2007), where hidden trading induces inefficiencies, Bisin and Rampini (2006) demonstrate that under limited government commitment, the potential for secret trades can enhance welfare by constraining the principal. While the First Welfare Theorem traditionally motivates markets on efficiency grounds, Bisin and Rampini (2006) highlight their role as commitment devices that address the planner's time inconsistency.

Brzustowski, Georgiadis-Harris, and Szentes (2023) introduce a hybrid model of limited commitment that complements the approaches discussed in this chapter. First, mechanism design is not restricted to one-period contracts—the principal can offer infinite-horizon dynamic *smart contracts*. Second, the principal retains the ability to propose new contracts in each period without requiring the agent's consent to discard prior agreements. The authors demonstrate that smart contracts can mitigate ratchet effects and, in particular, show that a durable-good seller can attain revenue above cost. However, we do not yet know what the canonical 'smart' mechanisms are in Brzustowski, Georgiadis-Harris, and Szentes (2023).

A recent literature studies dynamic contracting with a third party (a mediator) who facilitates the principal–agent relationship. Lomys and Yamashita (2022), building on Doval and Skreta (2022), analyze a long-lived *mediator* with memory who mediates communication between the principal and the agent. Insulated from dismissal, the mediator accumulates private information inaccessible to the principal. The authors show that such a mediator expands the set of implementable allocations by partially restoring commitment power. Trivially, if the mediator withholds all information from the designer, any outcome implementable under full commitment can also be achieved in this setting. However, even under natural assumptions about what the principal observes, Attar, Bozzoli, and Strausz (2025a) demonstrate that in certain dynamic contracting environments subject to renegotiation, dynamic communication mechanisms exist that uniquely implement the full-commitment allocation. Other forms of imperfect commitment and flexible contracting could arise as researchers investigate multi-agent settings—an area where, beyond Bester and Strausz (2000) and Skreta (2015), very little has been done.[26]

[26]Multi-agent settings also provide new opportunities for the designer to discipline herself: for example, if agents arrive sequentially, they can punish an opportunistic designer. This force is explored in Ekmekci et al. (2024), who consider a setting with a continuum of agents, and relates to the dynamic incentives in the literature on relational contracts stemming from the seminal contribution of Levin (2003). It is also somewhat reminiscent of the disciplining power of voters as a counterforce to the

References

Acemoglu, D., M. Golosov, and A. Tsyvinski (2010): "Dynamic Mirrlees taxation under political economy constraints," *The Review of Economic Studies*, 77, 841–881.

Acharya, A. and J. Ortner (2017): "Progressive learning," *Econometrica*, 85, 1965–1990.

Akbarpour, M. and S. Li (2020): "Credible auctions: A trilemma," *Econometrica*, 88, 425–467.

Attar, A., L. Bozzoli, and R. Strausz (2025a): "Self-Revealing Renegotiation," .

Attar, A., E. Campioni, T. Mariotti, and A. Pavan (2025b): "Keeping the Agents in the Dark: Competing Mechanisms, Private Disclosures, and the Revelation Principle," .

Attar, A., E. Campioni, T. Mariotti, and G. Piaser (2021): "Competing mechanisms and folk theorems: Two examples," *Games and Economic Behavior*, 125, 79–93.

Ausubel, L. M. and R. J. Deneckere (1989): "Reputation in bargaining and durable goods monopoly," *Econometrica*, 511–531.

Baron, D. P. and D. Besanko (1984): "Regulation and information in a continuing relationship," *Information Economics and policy*, 1, 267–302.

Battaglini, M. (2007): "Optimality and renegotiation in dynamic contracting," *Games and economic behavior*, 60, 213–246.

Beccuti, J. and M. Möller (2018): "Dynamic adverse selection with a patient seller," *Journal of Economic Theory*, 173, 95–117.

Bergemann, D. and J. Välimäki (2019): "Dynamic mechanism design: An introduction," *Journal of Economic Literature*, 57, 235–274.

Bester, H. and R. Strausz (2000): "Imperfect commitment and the revelation principle: the multi-agent case," *Economics Letters*, 69, 165–171.

——— (2001): "Contracting with imperfect commitment and the revelation principle: the single agent case," *Econometrica*, 69, 1077–1098.

——— (2007): "Contracting with imperfect commitment and noisy communication," *Journal of Economic Theory*, 136, 236–259.

Bisin, A. and A. A. Rampini (2006): "Markets as beneficial constraints on the government," *Journal of public Economics*, 90, 601–629.

planner's imperfect commitment explored in Acemoglu et al. (2010).

Breig, Z. (2022): "Repeated contracting without commitment," *Journal of Economic Theory*, 204, 105514.

Brzustowski, T., A. Georgiadis-Harris, and B. Szentes (2023): "Smart contracts and the coase conjecture," *American Economic Review*, 113, 1334–1359.

Bulow, J. I. (1982): "Durable-goods monopolists," *Journal of Political Economy*, 90, 314–332.

Calzolari, G. and A. Pavan (2006a): "Monopoly with resale," *The RAND Journal of Economics*, 37, 362–375.

——— (2006b): "On the optimality of privacy in sequential contracting," *Journal of Economic theory*, 130, 168–204.

Celik, G. (2015): "Implementation by gradual revelation," *The RAND Journal of Economics*, 46, 271–296.

Crawford, V. P. and J. Sobel (1982): "Strategic information transmission," *Econometrica*, 1431–1451.

Crémer, J. (1995): "Arm's Length Relationships," *The Quarterly Journal of Economics*, 110, 275–295.

Deb, R. and M. Said (2015): "Dynamic screening with limited commitment," *Journal of Economic Theory*, 159, 891–928.

Dewatripont, M. (1986): "On the Theory of Commitment with Applications to the Labor Market," Ph.D. thesis, Harvard University.

——— (1988): "Commitment through renegotiation-proof contracts with third parties," *The Review of Economic Studies*, 55, 377–390.

Dewatripont, M. and E. Maskin (1995): "Contractual Contingencies and Renegotiation," *The RAND Journal of Economics*, 26, 704–719.

Doval, L. and V. Skreta (2022): "Mechanism design with limited commitment," *Econometrica*, 90, 1463–1500.

——— (2024a): "Constrained information design," *Mathematics of Operations Research*, 49, 78–106.

——— (2024b): "Mechanism design with limited commitment: Markov environments," *European Economic Review*, 167, 104777.

——— (2025): "Purchase History and Product Personalization," *The RAND Journal of Economics*, n/a.

Dworczak, P. (2020): "Mechanism design with aftermarkets: Cutoff mechanisms," *Econometrica*, 88, 2629–2661.

Dworczak, P. and G. Martini (2019): "The simple economics of optimal persuasion," *Journal of Political Economy*, 127, 1993–2048.

Ekmekci, M., L. Maestri, and D. Wei (2024): "Turning the Ratchet: Dynamic Screening with Multiple Agents," *arXiv preprint arXiv:2405.04468*.

Farhi, E., C. Sleet, I. Werning, and S. Yeltekin (2012): "Non-linear capital taxation without commitment," *Review of Economic Studies*, 79, 1469–1493.

Fiocco, R. and R. Strausz (2015): "Consumer standards as a strategic device to mitigate ratchet effects in dynamic regulation," *Journal of Economics & Management Strategy*, 24, 550–569.

Freixas, X., R. Guesnerie, and J. Tirole (1985): "Planning under incomplete information and the ratchet effect," *The Review of Economic Studies*, 52, 173–191.

Fudenberg, D. and J. Tirole (1990): "Moral hazard and renegotiation in agency contracts," *Econometrica: Journal of the Econometric Society*, 1279–1319.

Gerardi, D. and L. Maestri (2020): "Dynamic contracting with limited commitment and the ratchet effect," *Theoretical Economics*, 15, 583–623.

Golosov, M. and L. Iovino (2021): "Social insurance, information revelation, and lack of commitment," *Journal of Political Economy*, 129, 2629–2665.

Golosov, M., V. Skreta, A. Tsyvinski, and A. Wilson (2014): "Dynamic strategic information transmission," *Journal of Economic Theory*, 151, 304–341.

Golosov, M. and A. Tsyvinski (2007): "Optimal taxation with endogenous insurance markets," *The Quarterly Journal of Economics*, 122, 487–534.

Golosov, M., A. Tsyvinski, I. Werning, P. Diamond, and K. L. Judd (2006): "New dynamic public finance: A user's guide [with comments and discussion]," *NBER macroeconomics annual*, 21, 317–387.

Gul, F., H. Sonnenschein, and R. Wilson (1986): "Foundations of dynamic monopoly and the coase conjecture," *Journal of Economic Theory*, 39, 155 – 190.

Hart, O. D. and J. Tirole (1988): "Contract renegotiation and Coasian dynamics," *The Review of Economic Studies*, 55, 509–540.

Hermann, E. and S. Puntoni (2024): "Artificial intelligence and consumer behavior: From predictive to generative AI," *Journal of Business Research*, 180, 114720.

Hiriart, Y., D. Martimort, and J. Pouyet (2011): "Weak enforcement of environmental policies: a tale of limited commitment and limited fines," *Annals of Economics and Statistics*, 25–42.

Kamenica, E. and M. Gentzkow (2011): "Bayesian persuasion," *American Economic Review*, 101, 2590–2615.

Kartik, N., S. Lee, and D. Rappoport (2024): "Single-crossing differences in convex environments," *Review of Economic Studies*, 91, 2981–3012.

Kolotilin, A. (2018): "Optimal information disclosure: A linear programming approach," *Theoretical Economics*, 13, 607–635.

Laffont, J.-J. (1994): "The new economics of regulation ten years after," *Econometrica: journal of the Econometric Society*, 507–537.

Laffont, J.-J. and J. Tirole (1988): "The dynamics of incentive contracts," *Econometrica*, 1153–1175.

——— (1993): *A theory of incentives in procurement and regulation*, MIT press.

Levin, J. (2003): "Relational incentive contracts," *American Economic Review*, 93, 835–857.

Liu, Q., K. Mierendorff, X. Shi, and W. Zhong (2019): "Auctions with limited commitment," *American Economic Review*, 109, 876–910.

Lomys, N. and T. Yamashita (2022): "A mediator approach to mechanism design with limited commitment," *Available at SSRN 4116543*.

McAdams, D. and M. Schwarz (2007): "Credible sales mechanisms and intermediaries," *American Economic Review*, 97, 260–276.

McAfee, R. P. and D. Vincent (1997): "Sequentially optimal auctions," *Games and Economic Behavior*, 18, 246–276.

Mussa, M. and S. Rosen (1978): "Monopoly and product quality," *Journal of Economic theory*, 18, 301–317.

Myerson, R. B. (1981): "Optimal auction design," *Mathematics of Operations Research*, 6, 58–73.

——— (1982): "Optimal coordination mechanisms in generalized principal–agent problems," *Journal of Mathematical Economics*, 10, 67–81.

——— (1986): "Multistage games with communication," *Econometrica: Journal of the Econometric Society*, 323–358.

Pavan, A. and G. Calzolari (2009): "Sequential contracting with multiple principals," *Journal of Economic Theory*, 144, 503–531.

Pavan, A., I. Segal, and J. Toikka (2014): "Dynamic mechanism design: A myersonian approach," *Econometrica*, 82, 601–653.

Rey, P. and B. Salanie (1996): "On the value of commitment with asymmetric information," *Econometrica: Journal of the Econometric Society*, 1395–1414.

Riley, J. and R. Zeckhauser (1983): "Optimal selling strategies: When to haggle, when to hold firm," *The Quarterly Journal of Economics*, 98, 267–289.

Roberts, K. W. (1984): "The theoretical limits to redistribution," *The Review of Economic Studies*, 51, 177–195.

Rockafellar, R. T. (1970): *Convex analysis*, Princeton University Press.

Salanié, B. (2005): *The economics of contracts: a primer*, MIT press.

Skreta, V. (2006): "Sequentially optimal mechanisms," *The Review of Economic Studies*, 73, 1085–1111.

——— (2015): "Optimal auction design under non-commitment," *Journal of Economic Theory*, 159, 854–890.

Sleet, C. and Ş. Yeltekin (2006): "Credibility and endogenous societal discounting," *Review of Economic Dynamics*, 9, 410–437.

——— (2008): "Politically credible social insurance," *Journal of monetary Economics*, 55, 129–151.

Stantcheva, S. (2020): "Dynamic taxation," *Annual Review of Economics*, 12, 801–831.

Stokey, N. L. (1981): "Rational expectations and durable goods pricing," *The Bell Journal of Economics*, 112–128.

Strausz, R. (2003): "Deterministic mechanisms and the revelation principle," *Economics Letters*, 79, 333–337.

Strotz, R. H. (1955): "Myopia and inconsistency in dynamic utility maximization," *The Review of Economic Studies*, 23, 165–180.

Sugaya, T. and A. Wolitzky (2021): "The revelation principle in multistage games," *The Review of Economic Studies*, 88, 1503–1540.

Taneva, I. (2019): "Information design," *American Economic Journal: Microeconomics*, 11, 151–185.

Townsend, R. M. (1988): "Information constrained insurance: the revelation principle extended," *Journal of Monetary Economics*, 21, 411–450.

Vartiainen, H. (2013): "Auction design without commitment," *Journal of the European Economic Association*, 11, 316–342.