EFFICIENT DYNAMIC MECHANISM DESIGN WITHOUT A COMMON PRIOR

By

Dirk Bergemann and Marek Bojko

September 2026

COWLES FOUNDATION DISCUSSION PAPER NO. 2540R1

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

http://cowles.yale.edu/

Efficient Dynamic Mechanism Design without a Common Prior*

Dirk Bergemann† Marek Bojko‡

September 1, 2026

Abstract

We study efficient dynamic mechanism design with independent private values when agents do not share a common prior over the stochastic environment. Each agent privately observes the kernel governing her own type evolution and may hold arbitrary beliefs about others' kernels. We include kernels in agents' type spaces and show that the dynamic team mechanism of Athey and Segal (2013) and the dynamic pivot mechanism of Bergemann and Välimäki (2010) implement the socially efficient allocation in periodic ex-post equilibrium. We further show that kernels need be elicited only at the outset and that these mechanisms induce the efficient private acquisition of kernels.

Keywords: Dynamic mechanism design, robust mechanism design, subjective priors, periodic ex-post equilibrium, dynamic pivot mechanism, dynamic team mechanism, VCG, supply-chain coordination.

JEL Classification: C73, D44, D47, D82, D83, D86.

*We are grateful to Jiaxuan Li, Garrett van Ryzin, Eva Tardos, and Elhanan Mossel for helpful conversations. We gratefully acknowledge financial support through the Office of Naval Research MURI grant N00014-23-1-2700 and the National Science Foundation through grants SES-2049754 and SES-2419015. Claude, ChatGPT, and Refine.ink were used for proofreading and basic text editing. All errors are our own.

†Department of Economics, Yale University, dirk.bergemann@yale.edu

‡Department of Economics, Yale University, marek.bojko@yale.edu

1 Introduction

When a retailer and a vendor make sequential decisions about pricing, production, and shipment in a supply chain, the efficiency of their coordination depends on how information is shared and how incentives are aligned. Standard models assume that all parties share a common prior over payoff-relevant states and their evolution. While analytically convenient, this assumption rarely holds in practice. A retailer forecasts demand using proprietary point-of-sale, search, and customer data, while a vendor assesses component availability using supplier relationships. There is little reason to expect these diverse sources of information to be derived from any common probabilistic model. More broadly, in dynamic procurement, regulation (Laffont and Tirole, 1993; Acemoglu and Lensman, 2024; Koh and Sanguanmoo, 2025), and public-good provision, the stochastic models governing payoffs are typically private to the agents who hold them.

This paper asks which parts of the common prior assumption are required for efficient dynamic mechanism design and which can be dispensed with. We give a precise answer for environments with independent private values. In our model, each agent privately knows the stochastic kernel governing the evolution of her own payoff type. We do not require the other agents, or the designer, to know that kernel. What is commonly known is the structure and interpretation of the stochastic environment: the set of admissible stochastic kernels for each agent, the persistence of each agent's kernel, and the independence of type transitions across agents. We impose no prior over which kernel profile obtains and we specify no beliefs or higher-order beliefs about them. We consider direct mechanisms that elicit both the stochastic kernels and the agents' realized payoff types, treating each agent's kernel as part of her type.

We show that two canonical mechanisms, the dynamic team mechanism of Athey and Segal (2013) and the dynamic pivot mechanism of Bergemann and Välimäki (2010), implement the socially efficient allocation in periodic ex-post equilibrium on this enlarged type space (Propositions 1 and 2, respectively). The key observation is that when the full type is elicited in each period, an agent's report of her kernel affects her payoff only through the allocation chosen by the mechanism, while her continuation payoffs are evaluated under her true kernel. Truthful reporting therefore aligns the agent's incentives with the planner's dynamic social surplus maximization problem. As we make precise in Section 4.1, both propositions can also be obtained from the implementation results of Athey and Segal (2013) and Bergemann and Välimäki (2010) applied to a suitably expanded type space and transition law. That this reduction is available is precisely the point: periodic ex-post implementation never invokes a prior over the (expanded) state, so the common prior over the stochastic environment does no work in these arguments. The contribution of this paper is to delineate precisely the informational assumptions behind these results: the mechanisms require a common language for describing the environment, but not common probabilistic assessments within that language.

As in static settings, a common prior coordinates beliefs about the realized private information of the other agents. In dynamic settings, it additionally specifies the stochastic law of motion under which continuation values are computed. The robust mechanism design literature in the tradition of the Wilson doctrine shows that the first role can be dispensed with in static settings by moving to ex-post or dominant-strategy solution concepts (Wilson, 1987; Bergemann and Morris, 2005; Chung and Ely, 2007). Our results show that, under independent private values, the second and distinctly dynamic role is also dispensable: the stochastic environment itself can be privately known and elicited by the mechanism.

We then use this observation in two ways. First, Proposition 3 and Corollary 1 show that, when kernels are persistent, the mechanisms can elicit kernels once at the outset and then ask agents only for their realized payoff types. Second, Proposition 4 shows that the same mechanisms induce efficient investment in kernels; for example, a vendor may invest in a production process that changes its failure rates.

Related literature. This work is related to three literatures. The first is the literature on efficient dynamic mechanism design, beginning with Baron and Besanko (1984) and including Bergemann and Välimäki (2010), Athey and Segal (2013), Pavan et al. (2014), and Csóka et al. (2023); see Bergemann and Välimäki (2019) for a survey and Bergemann et al. (2026) for an application to supply chain coordination. We clarify the role of the common prior assumption that underlies much of the existing literature.

The second is the literature on robust mechanism design, originating with Wilson (1987), including the belief-free approach developed by Bergemann and Morris (2005) in static environments and extended by Penta (2015) to dynamic environments. Different notions of robustness in static settings are studied by Chung and Ely (2007), Carroll (2017, 2019), and Koçyiğit et al. (2020); in dynamic settings, they are studied by, among others, Chassang (2013), Libgober and Mu (2021), Koh and Sanguanmoo (2025), and Csóka (2025). Our focus is distinct but complementary: we study how efficient mechanisms can elicit privately known stochastic kernels.

The third literature is the recent work on mechanism design under adversarial or learning-theoretic assumptions, which also moves beyond the common prior framework; see, for example, Camara et al. (2020) and Han et al. (2025). In contrast, we keep the Bayesian framework and ask how far the common prior assumption can be relaxed within it.

Outline. Section 2 introduces the model. Section 3 establishes the main results for the dynamic team and dynamic pivot mechanisms. Section 4 develops further interpretations and presents results on one-time kernel elicitation and efficient investment. Section 5 concludes.

2 Model

2.1 Setup

We consider a discounted discrete-time model with private values and quasi-linear utilities, following Athey and Segal (2013) and Bergemann and Välimäki (2010). There is a finite set of agents $i \in N \equiv \{1, \dots, n\}$ interacting in each period $t \in \{0, 1, \dots\}$. Each agent i in period t receives a payoff that depends on the current allocation $x_t \in X$, the current monetary transfer $p_{it} \in \mathbb{R}$, and the private *payoff type* $\theta_{it} \in \Theta_i$. We denote by $\Theta \equiv \bigtimes_{i=1}^n \Theta_i$ the product type space. The flow utility of agent i in period t takes the quasi-linear form

$$u_i(x_t, p_{it}, \theta_{it}) \equiv v_i(x_t, \theta_{it}) - p_{it}.$$

Agents share a discount factor $\delta \in (0, 1)$.

The payoff type θ_{it} of agent i follows a controlled Markov process on the state space Θ_i. The current type θ_{it} and the current allocation x_t determine the conditional distribution of the next-period type $\theta_{i,t+1}$ through a stochastic transition kernel

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

Let $\mathcal{F}_i$ denote the set of admissible kernels for agent i and let $\mathcal{F} \equiv \bigtimes_{i=1}^n \mathcal{F}_i$ denote the product space. We denote a kernel profile by $F \in \mathcal{F}$. The sets $\mathcal{F}_i$, $i \in N$, are common knowledge; they constitute the shared language in which the possible stochastic environments are described. Which kernel in $\mathcal{F}_i$ obtains is the private information of agent i. Conditional on the current state and allocation, agents' type transitions are independent across agents, so the aggregate transition kernel is

$$\prod_{i \in N} F_i(d\theta_{i,t+1} \mid \theta_{it}, x_t).$$

We impose the following regularity conditions throughout the paper.

Assumption 1 (Regularity)**.** (i) The allocation space X and each type space Θ_i are compact metric spaces, endowed with their Borel σ-algebras; (ii) each valuation function $v_i : X \times \Theta_i \to \mathbb{R}$ is continuous; we let $K > 0$ denote a common bound, $\|v_i\|_\infty \le K$ for every i; (iii) each $\mathcal{F}_i$ is a finite set of Borel transition kernels, and each $F_i \in \mathcal{F}_i$ is Feller continuous in (θ_i, x).

The finiteness of the kernel sets in Assumption 1(iii) makes the product space $\Theta_i \times \mathcal{F}_i$ a compact metric space (with $\mathcal{F}_i$ carrying the discrete topology), so that all objects defined on the expanded type space below are measurable without further qualification, and it lets us focus on the informational content of the results.

The transition kernel $F_i \in \mathcal{F}_i$ is private information of agent i, observed at the outset of period 0. At the beginning of each period t, the agent privately observes her realized payoff type θ_{it}. The private information of agent i at the beginning of period t is therefore the pair

$$\omega_{it} = \big(\theta_{it}, F_i\big) \in \Omega_i \equiv \Theta_i \times \mathcal{F}_i.$$

We denote by $\Omega \equiv \bigtimes_{i=1}^n \Omega_i$ the product space, with typical element $\omega_t = (\theta_t, F)$.

We emphasize two features of this formulation. First, no common prior over the kernel profile F is imposed: neither the modeler nor the designer specifies any joint distribution on $\mathcal{F}$, nor does any agent. Second, we do not specify agents' beliefs about other agents' kernels, or higher-order beliefs. Under the equilibrium concept introduced below, such beliefs play no role.

2.2 Social Efficiency

We focus on the utilitarian socially efficient allocation. Starting in any period t at state $\omega_t = (\theta_t, F)$, the planner solves

$$W(\omega_t) \equiv \max_{\{x_s\}_{s\geq t}} \mathbb{E}_F \left[\sum_{s=t}^{\infty} \delta^{s-t} \sum_{i=1}^{n} v_i(x_s(\omega_s), \theta_{is}) \mid \theta_t \right],$$

where the expectation is taken under the kernel profile F. Under Assumption 1, the planner's Bellman operator maps the space of bounded continuous functions on Ω into itself and is a contraction with modulus δ; its unique fixed point W is therefore bounded and continuous, and a measurable optimal policy exists by the measurable maximum theorem. Stated recursively, W solves the Bellman equation

$$W(\omega_t) = \max_{x_t \in X} \left\{ \sum_{i=1}^{n} v_i(x_t, \theta_{it}) + \delta\, \mathbb{E}_F \left[W(\omega_{t+1}) \,\middle|\, \theta_t, x_t \right] \right\}.$$

Denote by $\mathbf{x}^* = \{x_t^*\}_{t=0}^{\infty}$ the (measurable) socially optimal policy and by W_{-i} the analogous value function for the economy without agent i, evaluated under the kernel profile F_{-i}.

2.3 Dynamic Mechanisms and Implementation

We focus on direct mechanisms that implement the efficient policy in periodic ex-post equilibrium. We first restrict attention to mechanisms that elicit the full type ω_{it} from each agent i in every period t; Section 4.2 shows that the same allocation can be implemented through a one-time elicitation of kernels at the outset.

Denote a report of agent i in period t by $r_{it} \in \Omega_i$ and a profile of period-t reports by $r_t = (\hat{\theta}_t, \hat{F}) \in \Omega$. We assume that the sequence of past reports is public information. The public history at the beginning of period t,

$$h_t = (r_0, x_0, r_1, x_1, \ldots, r_{t-1}, x_{t-1}),$$

collects past reports and allocations; the set of feasible public histories is denoted H_t, with $h_0 = \emptyset$. The private history of agent i in period t further records the agent's realized types,

$$h_{it} = (\omega_{i0}, r_0, x_0, \omega_{i1}, r_1, x_1, \ldots, \omega_{i,t-1}, r_{t-1}, x_{t-1}, \omega_{it}).$$

The set of such private histories is denoted H_{it}.

An efficient dynamic direct mechanism is a tuple $(\mathbf{x}^*, \mathbf{p})$, where $\mathbf{x}^*$ and $\mathbf{p}$ are sequences of allocation and transfer functions

$$x_t^*:\Omega\to X,\qquad p_t:\Omega\times H_t\to\mathbb{R}^n,$$

and where $\mathbf{x}^*$ is the socially efficient allocation policy. The transfer p_t may depend on the full public history. Throughout, the allocation and transfer functions are evaluated at the *reported* profile r_t; in particular, wherever a conditional expectation entering a transfer formula carries a kernel index, that kernel is the reported one, so that the mechanism can compute the transfer from the reports alone. Expectations in the agents' payoff calculations are, by contrast, always taken under the true kernels.

We call a mechanism $(\mathbf{x}^*,\mathbf{p})$ *admissible* if each p_t is Borel measurable and transfers are uniformly bounded: there exists $\bar{P}>0$ such that $|p_{it}(r_t,h_t)|\le\bar{P}$ for every i, t, r_t, and h_t. Admissibility guarantees that expected discounted payoffs are well defined and that the dynamic-programming arguments below apply. We restrict attention to admissible mechanisms throughout. Both the dynamic team mechanism and the dynamic pivot mechanism are admissible.

A pure reporting strategy of agent i is a measurable mapping $r_{it}:H_{it}\to\Omega_i$. Fixing a mechanism $(\mathbf{x}^*,\mathbf{p})$, the expected discounted payoff of agent i from the strategy $\mathbf{r}_i=\{r_{it}\}_{t=0}^\infty$, when other agents use strategies $\mathbf{r}_{-i}$ and types are drawn from the true kernel profile F, is

$$\mathbb{E}_F\left[\sum_{t=0}^{\infty}\delta^t\big[\,v_i(x_t^*(r_t),\theta_{it})-p_{it}(r_t,h_t)\,\big]\,\Big|\,\theta_0\right].$$

When the other agents report truthfully, we denote this payoff by $U_i(\omega_0;\mathbf{r}_i)$. Still fixing truthful reporting by the other agents, the agent's continuation value satisfies the recursion

$$V_i^*(\omega_t,h_t)=\sup_{r_{it}\in\Omega_i}\Big\{\,v_i(x_t^*(r_{it},\omega_{-it}),\theta_{it})-p_{it}\big((r_{it},\omega_{-it}),h_t\big)+\delta\;\mathbb{E}_F\big[\,V_i^*(\omega_{t+1},h_{t+1})\,\big|\,\theta_t,x_t^*(r_{it},\omega_{-it})\,\big]\Big\}.$$

Because flow utilities are uniformly bounded under an admissible mechanism, the operator defined by the right-hand side is a contraction with modulus δ on the space of bounded functions of (ω_t,h_t), and V_i^* is its unique bounded fixed point.[1] The value function V_i^* represents the continuation value of the agent given public history h_t and the current type profile ω_t, when the rest of the agents report truthfully. Our equilibrium concept is periodic ex-post equilibrium.[2]

Definition 1 (Periodic ex-post incentive compatibility)**.** The dynamic direct mechanism $(\mathbf{x}^*,\mathbf{p})$ is *periodic ex-post incentive compatible* if, for every agent i, every public history $h_t\in H_t$, and every

[1]For uncountable Θ_i, the measurability questions raised by the supremum are handled by the standard universally measurable dynamic-programming framework of Bertsekas and Shreve (1978).

[2]For periodic ex-post incentive compatible mechanisms, as defined next, Lemma 1 in Appendix A verifies that V_i^* coincides with the supremum of the agent's expected discounted payoffs over reporting strategies.

type profile $\omega_t \in \Omega$, truthful reporting is a best response when all other agents report truthfully:

$$\omega_{it} \in \arg\max_{r_{it}\in\Omega_i}\Big\{ v_i\big(x_t^*(r_{it},\omega_{-it}),\theta_{it}\big) - p_{it}\big((r_{it},\omega_{-it}),h_t\big) + \delta\, \mathbb{E}_F\big[\, V_i^*(\omega_{t+1},h_{t+1}) \,\big|\, \theta_t, x_t^*(r_{it},\omega_{-it}) \,\big] \Big\}.$$

Periodic ex-post incentive compatibility requires truth-telling to be optimal state by state: the displayed condition holds for *every* realization ω_{-it} of the contemporaneous types of the other agents, not merely in expectation over them. The mechanism discloses nothing about ω_{-it} to agent i before she reports; the quantifier over ω_{-it} is a robustness requirement on the mechanism, not an assumption that agents observe each other's types. Because the contemporaneous types (θ_{-it}, F_{-i}) of the other agents exhaust all payoff-relevant information about the current and future evolution of their payoff types, and because the inequality holds at every realization of these types, truthful reporting is optimal under *any* belief that agent i might hold about the other agents' payoff types or kernels. Hence, we need not specify agent i's beliefs over the types of the other agents, including F_{-i}, for any agent i.

3 Socially Efficient Dynamic Mechanisms

In this section, we show that the dynamic team mechanism of Athey and Segal (2013) and the dynamic pivot mechanism of Bergemann and Välimäki (2010) implement the socially efficient allocation on the expanded type space. In both mechanisms, periodic ex-post incentive compatibility follows because each agent is made a residual claimant of the full social surplus. As noted in Section 4.1, both results can be obtained directly from the corresponding implementation results of Athey and Segal (2013) and Bergemann and Välimäki (2010) by suitably expanding the type space. For expositional clarity, the arguments below explicitly distinguish between the elicitation of payoff types and the elicitation of stochastic kernels, thereby clarifying the role played by each component of the agent's type.

3.1 Dynamic Team Mechanism

In the *dynamic team mechanism*, we set for each agent i and every reported profile $r_t = (\hat\theta_t, \hat F) \in \Omega$,

$$p_{it}(r_t) = -\sum_{j\neq i} v_j\big(x_t^*(r_t), \hat\theta_{jt}\big),$$

independently of the public history. In words, each agent is paid in each period the gross flow surplus that the other agents receive at the efficient allocation, evaluated at the current reports. Note that the reported kernels $\hat F$ enter the transfer only through the induced allocation $x_t^*(r_t)$.

Proposition 1 (Dynamic team mechanism). *The dynamic team mechanism is periodic ex-post incentive compatible.*

Proof. Flow utilities under the team mechanism are uniformly bounded, so discounted payoffs are continuous at infinity and the one-stage deviation principle applies. It therefore suffices to consider a deviation by agent i in a single period t, reporting $r_{it} \neq \omega_{it}$ while reporting truthfully thereafter, against truthful reports of the others. Under the team transfer, the agent's continuation payoff from such a deviation is

$$\sum_{j=1}^{n} v_j\big(x_t^*(r_{it}, \omega_{-it}), \theta_{jt}\big) + \delta\, \mathbb{E}_F[W(\omega_{t+1}) \mid \theta_t, x_t^*(r_{it}, \omega_{-it})]\,,$$

where the expectation is taken under the true kernel profile F: agent i's own type evolves according to her true kernel regardless of her report, and the other agents report truthfully. That is, agent i's objective coincides with the planner's Bellman objective evaluated at the allocation induced by the report r_{it}. Since x_t^* is chosen to maximize this expression, the agent's payoff is maximized by reporting ω_{it} truthfully. Hence no profitable one-stage deviation exists. □

A report of F_i affects agent i's payoff only through the induced allocation path, while her continuation expectations are evaluated under her true kernel. The agent's incentives align with the planner's because reports affect allocations, while continuation payoffs are evaluated under the true transition law.

3.2 Dynamic Pivot Mechanism

The *dynamic pivot mechanism* makes each agent's payoff equal to her dynamic marginal contribution to social welfare. At state ω_t, the dynamic marginal contribution of agent i is

$$M_i(\omega_t) \equiv W(\omega_t) - W_{-i}(\omega_{-i,t}),$$

where W is the efficient value function in the full economy and W_{-i} is the efficient value function in the economy without agent i. The corresponding *flow marginal contribution* m_i is defined by

$$M_i(\omega_t) = m_i(\omega_t) + \delta\, \mathbb{E}_F[M_i(\omega_{t+1}) \mid \theta_t, x_t^*(\omega_t)]\,.$$

Equivalently, Bergemann and Välimäki (2010) derive the following expression using the fact that, by conditional independence across agents, the transition of $\omega_{-i,t+1}$ depends only on $(\theta_{-i,t}, x_t)$ and F_{-i},

$$\begin{aligned} m_i(\omega_t) = \sum_{j=1}^{n} v_j\big(x_t^*(\omega_t), \theta_{jt}\big) - \sum_{j\neq i} v_j\big(x_{-i,t}^*(\omega_{-i,t}), \theta_{jt}\big) \\ + \delta\left(\mathbb{E}_{F_{-i}}[W_{-i}(\omega_{-i,t+1}) \mid \theta_{-i,t}, x_t^*(\omega_t)] - \mathbb{E}_{F_{-i}}\left[W_{-i}(\omega_{-i,t+1}) \mid \theta_{-i,t}, x_{-i,t}^*(\omega_{-i,t})\right]\right), \end{aligned}$$

where $x_{-i,t}^*$ denotes the socially efficient allocation in the economy without agent i.

The pivot transfer is defined for every reported profile $r_t = (\hat{\theta}_t, \hat{F}) \in \Omega$ by evaluating these formulas at the current reports:

$$p_{it}(r_t) = v_i\big(x_t^*(r_t), \hat{\theta}_{it}\big) - m_i(r_t), \tag{1}$$

where $m_i(r_t)$ is the flow marginal contribution formula above evaluated at $(\hat{\theta}_t, \hat{F})$; in particular, the two conditional expectations in $m_i(r_t)$ are computed under the *reported* kernels $\hat{F}_{-i}$ of the other agents, and W_{-i} is the value function associated with the reported profile $\hat{F}_{-i}$. The transfer in (1) is thus well defined for every report profile, on and off the equilibrium path, and is computed by the mechanism from the current reports alone. Note that agent i's own reported kernel $\hat{F}_i$ enters her transfer only through the current allocation $x_t^*(r_t)$: the expectations in $m_i(r_t)$ involve only the reported kernels of the other agents. Under the transfer (1), agent i's flow utility at truthful reports equals her flow marginal contribution, so the agent internalizes her dynamic externality on the rest of society.

Proposition 2 (Dynamic pivot mechanism)**.** *The dynamic pivot mechanism is periodic ex-post incentive compatible.*

Proof. Flow utilities under the pivot mechanism are uniformly bounded, so payoffs are continuous at infinity and the principle of optimality applies: it suffices to show that if agent i receives as continuation value her marginal contribution, then truth-telling is optimal for i in period t:

$$\begin{aligned} &v_i(x_t^*(\omega_t), \theta_{it}) - p_{it}(\omega_t) + \delta\, \mathbb{E}_F[M_i(\omega_{t+1}) \mid \theta_t, x_t^*(\omega_t)] \\ &\qquad\qquad \geq v_i(x_t^*(r_{it}, \omega_{-i,t}), \theta_{it}) - p_{it}(r_{it}, \omega_{-i,t}) + \delta\, \mathbb{E}_F[M_i(\omega_{t+1}) \mid \theta_t, x_t^*(r_{it}, \omega_{-i,t})] \end{aligned}$$

for all $r_{it} \in \Omega_i$ and $\omega_t \in \Omega$; here the other agents report truthfully, so $\hat{F}_{-i} = F_{-i}$ and $\hat{\theta}_{-i,t} = \theta_{-i,t}$, while the expectations over ω_{t+1} are taken under the true kernel profile F. By construction of p_{it}, the left-hand side equals the marginal contribution $M_i(\omega_t)$ of agent i. Plugging in for M_i on the right-hand side, the inequality above is equivalent to

$$\begin{aligned} W(\omega_t) - W_{-i}(\omega_{-i,t}) &\geq v_i(x_t^*(r_{it}, \omega_{-i,t}), \theta_{it}) - p_{it}(r_{it}, \omega_{-i,t}) \\ &\quad + \delta\left(\mathbb{E}_F[W(\omega_{t+1}) \mid \theta_t, x_t^*(r_{it}, \omega_{-i,t})] - \mathbb{E}_{F_{-i}}[W_{-i}(\omega_{-i,t+1}) \mid \theta_{-i,t}, x_t^*(r_{it}, \omega_{-i,t})]\right). \end{aligned}$$

Further, writing out the transfer (1) at the report profile $(r_{it}, \omega_{-i,t})$ and canceling the valuation terms of agent i, we obtain

$$\begin{aligned} p_{it}(r_{it}, \omega_{-i,t}) &= \sum_{j \neq i} \left[v_j(x_{-i,t}^*(\omega_{-i,t}), \theta_{jt}) - v_j(x_t^*(r_{it}, \omega_{-i,t}), \theta_{jt})\right] \\ &\quad + \delta\left(\mathbb{E}_{F_{-i}}[W_{-i}(\omega_{-i,t+1}) \mid \theta_{-i,t}, x_{-i,t}^*(\omega_{-i,t})] - \mathbb{E}_{F_{-i}}[W_{-i}(\omega_{-i,t+1}) \mid \theta_{-i,t}, x_t^*(r_{it}, \omega_{-i,t})]\right). \end{aligned}$$

In particular, the transfer does not depend on the reported kernel $\hat{F}_i$ or the reported payoff type $\hat{\theta}_{it}$, except through the induced current allocation $x_t^*(r_{it}, \omega_{-i,t})$. Plugging in for $p_{it}(r_{it}, \omega_{-i,t})$, collecting terms, and using the Bellman equation for the economy without agent i, $W_{-i}(\omega_{-i,t}) = \sum_{j\neq i} v_j(x^*_{-i,t}(\omega_{-i,t}), \theta_{jt}) + \delta\, \mathbb{E}_{F_{-i}}[W_{-i}(\omega_{-i,t+1}) \mid \theta_{-i,t}, x^*_{-i,t}(\omega_{-i,t})]$, the inequality reduces to

$$W(\omega_t) - W_{-i}(\omega_{-i,t}) \geq \sum_{j=1}^{n} v_j(x^*_t(r_{it}, \omega_{-i,t}), \theta_{jt}) + \delta\, \mathbb{E}_F[W(\omega_{t+1}) \mid \theta_t, x^*_t(r_{it}, \omega_{-i,t})] - W_{-i}(\omega_{-i,t}).$$

This holds for all $r_{it} \in \Omega_i$ by the optimality of $\mathbf{x}^*$, completing the proof. □

The transfers of both mechanisms depend on the kernels reported by the other agents, but not directly on agent i's own reported kernel. In the truthful equilibrium, the kernels reported by the other agents coincide with their true kernels, so the stochastic model embedded in agent i's transfer is the true model governing the other agents' transitions.

4 Discussion and Extensions

4.1 Expanded Type Spaces and the Role of a Common Prior

So far, we have explicitly distinguished the two components of each agent's private information: the realized payoff type θ_{it} and the stochastic kernel F_i. This distinction is useful because the two objects play different roles. The payoff type determines the agent's current flow payoff, whereas the kernel governs the evolution of the agent's payoff type. In this section, we clarify how the current environment can be represented as a standard dynamic private-value environment on an expanded type space.

Let $\omega_i = (\theta_i, F_i) \in \Omega_i$. For $\omega_i = (\theta_i, F_i)$ and $\omega_i' = (\theta_i', F_i')$, define agent i's transition kernel on the expanded type space by

$$\mu_i(d\theta_i', dF_i' \mid \theta_i, F_i, x) = F_i(d\theta_i' \mid \theta_i, x)\, \delta_{F_i}(dF_i'),$$

where δ_{F_i} denotes the Dirac measure at F_i. Thus, the payoff component evolves according to the kernel contained in the agent's current type, while the kernel component remains fixed. Conditional independence across agents gives the aggregate transition kernel

$$\mu(d\omega' \mid \omega, x) = \prod_{i\in N} \mu_i(d\omega_i' \mid \omega_i, x).$$

The mapping μ is a commonly known transition rule on the expanded type space Ω. We have thus "markovized" the environment (Athey and Segal, 2013).

The relevant incentive arguments of Athey and Segal (2013) and Bergemann and Välimäki (2010) can then be applied on the expanded type space. The dynamic team mechanism makes each agent a residual claimant for total surplus, whereas the dynamic pivot mechanism assigns each agent her dynamic marginal contribution. Propositions 1 and 2 may therefore be viewed as direct applications of these constructions after replacing the payoff type θ_i with the expanded type (θ_i, F_i).

Taken together, the preceding section and the discussion above clarify the role of the common prior assumption. What is commonly known is the structure and interpretation of the environment: the sets $\mathcal{F}_i$ of admissible kernels, the persistence of each agent's kernel, the conditional independence of payoff-type transitions across agents, and the fact that F_i governs agent i's transition law. Each agent knows her own realized kernel, but neither the designer nor the other agents need know it. Moreover, the model imposes no initial probability distribution over the expanded state ω_0, and hence no common prior over the initial payoff types or the realized kernel profile. Nor do the results restrict agents' beliefs or higher-order beliefs about the kernels and payoff types of others. They therefore apply across compatible belief hierarchies, including hierarchies that do not admit a common prior. The contribution is thus to identify precisely the informational assumptions required by these mechanisms: a common description of the conditional evolution of the expanded state is sufficient; a common prior over its initial realization is not necessary. A commonly known stochastic kernel for types is without loss once the type space is expanded sufficiently. Conversely, nothing beyond the expanded type space and its commonly known transition structure is needed.

Nevertheless, explicitly distinguishing payoff types from stochastic kernels is conceptually useful. The distinction also matters for the results below. Persistence permits the kernel to be elicited once, at the outset, while interpreting F_i as a transition technology makes it meaningful to study costly investments that affect it.

4.2 Eliciting Kernels at the Outset

In the baseline formulation, the mechanism asks each agent to report the pair (θ_{it}, F_i) in every period, even though the kernel F_i is fixed over time. We now show it is sufficient to elicit stochastic kernels only once, at the outset. The argument proceeds by reduction: any admissible mechanism that is periodic ex-post incentive compatible in the baseline setting induces a mechanism that elicits the kernels only in period 0 and is incentive compatible when payoffs are evaluated at the outset. We discuss the augmented setting here informally and provide the formal construction in Appendix A.

Fix an admissible direct mechanism $(\mathbf{x}^*, \mathbf{p})$ in the environment of Section 2. Its *outset version* asks agent i to report (θ_{i0}, F_i) in period 0 and only θ_{it} in each period $t \geq 1$. After period 0, the mechanism stores the kernel profile reported at the outset and, in every later period, evaluates the original allocation and transfer formulas as if that same kernel profile were reported again. Equivalently, the public history fed into the original mechanism is obtained from the reduced public history by re-attaching the period 0 kernel report to every later payoff-type report.

We say that the outset version is *ex-post incentive compatible at the outset* if, for every agent i, every realized initial state (θ_0, F), and every reporting strategy $\mathbf{r}_i$, reporting truthfully yields a weakly higher expected payoff than following $\mathbf{r}_i$, given truthful reporting by the other agents.

Proposition 3 (One-time elicitation of kernels)**.** *Suppose that the admissible direct mechanism* $(\mathbf{x}^*, \mathbf{p})$ *is periodic ex-post incentive compatible in the repeated-report environment. Then its outset* *version is ex-post incentive compatible at the outset.*

Any outset reporting strategy induces a feasible repeated-report strategy: use the same period-0 report, repeat its kernel component thereafter, and keep the same payoff-type reports. Since the outset version stores the period-0 kernel reports and uses them in the original allocation and transfer rules, the induced strategy generates the same allocation and transfer path against truthful opponents. Hence any profitable outset deviation would yield a profitable deviation in the repeated-report environment, contradicting periodic ex-post incentive compatibility. Reporting truthfully is therefore optimal. Appendix A gives the formal argument. Combining Propositions 1, 2, and 3 yields the following corollary.

Corollary 1 (Eliciting kernels at the outset in the dynamic team and pivot mechanisms)**.** *The outset versions of the dynamic team mechanism and the dynamic pivot mechanism are ex-post incentive compatible at the outset.*

4.3 Efficient Investment

We now add an investment stage in which agents may acquire their kernels before participating in the mechanism. The acquired kernel is the law that governs the subsequent evolution of the agent's payoff type. Acquisition is therefore an investment in the agent's transition technology; for example, a vendor may invest in a production process that changes its failure rates.

The timing is as follows. First, the initial payoff-type profile θ_0 is realized and is commonly observed by the agents; the designer need not observe it.[3] Second, agents simultaneously and privately acquire kernels: an acquisition strategy of agent i is a measurable map $a_i : \Theta \to \mathcal{F}_i$, and acquiring $F_i \in \mathcal{F}_i$ costs $c_i(F_i)$, where $c_i : \mathcal{F}_i \to \mathbb{R}_+$; the cost is paid before the dynamic mechanism begins, after which F_i becomes agent i's private kernel as in the baseline model. Third, the dynamic mechanism is played as before. Fix θ_0 and define the cost-adjusted social surplus

$$\Phi(\theta_0, F) \equiv W(\theta_0, F) - \sum_{i \in N} c_i(F_i).$$

The efficient acquisition problem is to maximize Φ over $F \in \mathcal{F}$; since $\mathcal{F}$ is finite, a solution $F^* = F^*(\theta_0)$ exists, and since $W(\cdot, F)$ is continuous for each of the finitely many $F \in \mathcal{F}$, the selection $\theta_0 \mapsto F^*(\theta_0)$ can be taken measurable. We say a mechanism implements efficient investment and allocation in *ex-post equilibrium at the outset* if, for every agent i, every realized θ_0, and every pair $(F_i, \mathbf{r}_i)$ consisting of an alternative kernel and a reporting strategy, acquiring $F_i^*(\theta_0)$ and reporting truthfully yields a weakly higher expected payoff than acquiring F_i and following $\mathbf{r_i}$, given that other agents acquire $F_{-i}^*(\theta_0)$ and report truthfully. Because θ_0 is commonly observed at the acquisition stage, the acquisition strategies $a_i(\theta_0) = F_i^*(\theta_0)$ are feasible.

[3]We proceed with this assumption because the efficient kernel profile $F^*(\theta_0)$ defined below may depend on the entire initial profile θ_0. Alternatively, we could specify a prior over θ_0 and analyze investment from an ex-ante perspective.

Proposition 4 (Costly acquisition of kernels)**.** *The dynamic team mechanism and the dynamic pivot mechanism implement efficient investment and allocation in ex-post equilibrium at the outset.*

We present the proof in Appendix B. The argument builds on ideas from Rogerson (1992), Bergemann and Välimäki (2002), and Athey and Segal (2007). In particular, Athey and Segal (2007) show that the dynamic team mechanism induces efficient investment in the stochastic evolution of agents' types in their environment with independent private values.

Both the dynamic team mechanism and the dynamic pivot mechanism align each agent's unilateral investment incentives with cost-adjusted social surplus. Specifically, for any $F_i, \widehat{F}_i \in \mathcal{F}_i$ and any fixed $F_{-i} \in \mathcal{F}_{-i}$, agent i's truthful expected payoff gain from choosing $\widehat{F}_i$ rather than F_i is

$$\Phi(\theta_0, \widehat{F}_i, F_{-i}) - \Phi(\theta_0, F_i, F_{-i}).$$

This mirrors the observation of Quint and Kojima (2026) in a static environment. It follows that the acquisition game induced by either mechanism is an exact potential game with potential $\Phi(\theta_0, \cdot)$ (Monderer and Shapley, 1996). Consequently, every global maximizer of Φ, in particular $F^*(\theta_0)$, constitutes a Nash equilibrium of the acquisition stage, provided agents report truthfully in the continuation game.[4]

4.4 Evolving Subjective Models

The formal analysis treats each kernel F_i as fixed, and we have so far treated the kernels as representing an objective feature of the environment. We briefly discuss an informal extension in which agents' kernels are subjective and agents may revise these models over time.

One possible interpretation is as follows. In each period, agents report their current kernels, and the planner solves the continuation problem under the reported kernel profile. If an agent subsequently revises her model, the planner solves the problem again using the newly reported profile. Each current allocation is then efficient relative to the models held at that time and under the maintained assumption that those models will persist. The realized allocation path need not be optimal under any single intertemporal criterion.

This interpretation is consistent with Propositions 1 and 2 only in a limited sense. Their incentive arguments evaluate continuation payoffs as if the currently reported kernels were permanent. They therefore apply at each date if agents treat their current models as their best permanent descriptions of the environment and do not anticipate subsequent revisions. If agents instead have probabilistic beliefs about how their models will change, the model state and its evolution must be incorporated into an expanded type space. Anticipated revisions not represented in the reported state fall outside our formal analysis. Likewise, evaluating welfare under changing models would require an explicit criterion for aggregating assessments across time and agents.

[4] The acquisition stage may possess further equilibria. Proposition 4 is a partial implementation result.

5 Concluding Remarks

Under independent private values, efficient dynamic implementation requires neither common knowledge of agents' beliefs nor common knowledge of which stochastic environment obtains. Each agent's kernel may instead be privately known, without specifying a prior over kernel profiles, beliefs about other agents' kernels, or higher-order beliefs. Within this structure, the dynamic team and dynamic pivot mechanisms implement the efficient allocation in periodic ex-post equilibrium; moreover, the kernels need be elicited only once, and the same mechanisms induce their efficient costly acquisition. The results do not dispense with common stochastic structure altogether, but separate common understanding of what a stochastic model means from common knowledge of which model governs the environment. This distinction is particularly relevant when stochastic models are proprietary, as in supply-chain coordination (Bergemann et al., 2026).

An important direction for future research is to weaken the informational requirements further. For example, Penta (2015) introduces a notion of belief-free dynamic implementation. In the present setting, however, if agents' beliefs are genuinely subjective and may differ across agents, then there is no canonical "belief-free" or "prior-free" efficient rule, except in the conditional sense described above: efficiency can only be evaluated relative to the stochastic environment or belief system under consideration. One may therefore need to turn to fully prior-free or adversarial environments of the kind often studied in the computer science literature (Camara et al., 2020; Cesa-Bianchi and Lugosi, 2006).

A Details from Section 4.2

Throughout this appendix, fix an admissible direct mechanism $(\mathbf{x}^*, \mathbf{p})$, an agent i, and suppose that the agents other than i report truthfully. Recall that flow utilities are then uniformly bounded by $B \equiv K + \bar{P}$, so that $|V_i^*| \leq B/(1-\delta)$ for the unique bounded fixed point V_i^* of the recursion in Section 2. We first present the following result verifying optimality of truthful reporting in periodic ex-post incentive compatible mechanisms.

Lemma 1. *Let $(\mathbf{x}^*, \mathbf{p})$ be admissible and periodic ex-post incentive compatible. Then, for every initial state $\omega_0 \in \Omega$: (i) for every reporting strategy $\mathbf{r}_i$, $U_i(\omega_0; \mathbf{r}_i) \leq V_i^*(\omega_0, h_0)$; and (ii) the truthful strategy attains this bound: $U_i(\omega_0; \mathbf{r}_i^{\mathrm{tr}}) = V_i^*(\omega_0, h_0)$.*

Proof. (i) By definition of V_i^* as the fixed point of the Bellman operator, for every (ω_t, h_t) and every report $r_{it} \in \Omega_i$,

$$V_i^*(\omega_t, h_t) \geq v_i\big(x_t^*(r_{it}, \omega_{-it}), \theta_{it}\big) - p_{it}\big((r_{it}, \omega_{-it}), h_t\big) + \delta\, \mathbb{E}_F\big[\, V_i^*(\omega_{t+1}, h_{t+1}) \,\big|\, \theta_t, x_t^*(r_{it}, \omega_{-it}) \,\big]. \quad (2)$$

Fix a strategy $\mathbf{r}_i$ and apply (2) at the report prescribed by $\mathbf{r}_i$ at each private history. Taking expectations under the true kernel profile F along the induced process and iterating from period 0 to period $T-1$ yields

$$V_i^*(\omega_0, h_0) \geq \mathbb{E}_F\left[\sum_{t=0}^{T-1} \delta^t u_{it} \,\middle|\, \theta_0\right] + \delta^T\, \mathbb{E}_F\big[V_i^*(\omega_T, h_T) \,\big|\, \theta_0\big],$$

where u_{it} denotes agent i's realized flow utility in period t under $(\mathbf{r}_i, \mathbf{r}_{-i}^{\mathrm{tr}})$. Since $|u_{it}| \leq B$ and $|V_i^*| \leq B/(1-\delta)$, the tail of the payoff series and the terminal term are each bounded in absolute value by $\delta^T B/(1-\delta)$. Hence

$$V_i^*(\omega_0, h_0) \geq U_i(\omega_0; \mathbf{r}_i) - \frac{2\delta^T B}{1-\delta} \qquad \text{for every } T,$$

and letting $T \to \infty$ gives $V_i^*(\omega_0, h_0) \geq U_i(\omega_0; \mathbf{r}_i)$.

(ii) Periodic ex-post incentive compatibility (Definition 1) states that at every (ω_t, h_t) the truthful report attains the supremum in the recursion, so (2) holds with equality along truthful play. Iterating the equality and passing to the limit exactly as in part (i) yields $U_i(\omega_0; \mathbf{r}_i^{\mathrm{tr}}) = V_i^*(\omega_0, h_0)$. □

We now provide the formal construction of the outset mechanism. We augment the message space of each agent as follows. Agent i sends a message $m_{i0} \in M_{i0} \equiv \Theta_i \times \mathcal{F}_i$ in period 0, and a message $m_{it} \in M_{it} \equiv \Theta_i$ in each period $t \geq 1$. Let $M_0 \equiv \Theta \times \mathcal{F}$ and $M_t \equiv \Theta$ for each $t \geq 1$. A public history at the beginning of period t is

$$\hat{h}_t = (m_0, x_0, m_1, x_1, \ldots, m_{t-1}, x_{t-1}),$$

where $m_0 = (\hat{\theta}_0, G) \in M_0$ and $m_s \in M_s$ for each $s \geq 1$. Given a reduced public history $\hat{h}_t$, define its *lift* to the repeated-report environment by

$$\Lambda_t(\hat{h}_t) \equiv \big((\hat{\theta}_0, G), x_0, (m_1, G), x_1, \dots, (m_{t-1}, G), x_{t-1}\big) \in H_t.$$

That is, $\Lambda_t(\hat{h}_t)$ is obtained by attaching the period 0 kernel report G to every subsequent payoff-type report.

The *outset mechanism* associated with $(\mathbf{x}^*, \mathbf{p})$ is the mechanism $(\hat{\mathbf{x}}^*, \hat{\mathbf{p}})$ defined by

$$\hat{x}_0^*(m_0) \equiv x_0^*(m_0) \quad \text{and} \quad \hat{p}_0(m_0) \equiv p_0(m_0, \emptyset)$$

in period 0, and, for each $t \geq 1$,

$$\hat{x}_t^*(m_t, m_0) \equiv x_t^*\big((m_t, G)\big) \quad \text{and} \quad \hat{p}_t(m_t, \hat{h}_t) \equiv p_t\big((m_t, G), \Lambda_t(\hat{h}_t)\big).$$

Hence the outset mechanism stores the kernel profile reported in period 0 and, from period 1 onward, evaluates the original allocation and transfer formulas using that stored profile. Note that the outset mechanism inherits admissibility from $(\mathbf{x}^*, \mathbf{p})$.

A pure reporting strategy for agent i in the outset mechanism is a sequence $\hat{\mathbf{r}}_i = (\hat{r}_{i0}, \hat{r}_{i1}, \dots)$, where $\hat{r}_{i0} : \Theta_i \times \mathcal{F}_i \to \Theta_i \times \mathcal{F}_i$ and, for each $t \geq 1$, $\hat{r}_{it} : \hat{H}_{it} \to \Theta_i$, with $\hat{H}_{it}$ denoting agent i's private history in the reduced-message environment. Let $\hat{\mathbf{r}}_i^{\mathrm{tr}}$ denote the truthful strategy, defined by

$$\hat{r}_{i0}^{\mathrm{tr}}(\theta_{i0}, F_i) \equiv (\theta_{i0}, F_i) \quad \text{and} \quad \hat{r}_{it}^{\mathrm{tr}}(\hat{h}_{it}) \equiv \theta_{it}, \qquad t \geq 1,$$

where θ_{it} is the current payoff type contained in $\hat{h}_{it}$.

For every realized initial state (θ_0, F) and every strategy profile in which agents other than i are truthful, let $\hat{U}_i(\theta_0, F; \hat{\mathbf{r}}_i)$ denote agent i's expected discounted utility in the outset mechanism. We say that $(\hat{\mathbf{x}}^*, \hat{\mathbf{p}})$ is *ex-post incentive compatible at the outset* if, for every agent i, every realized initial state (θ_0, F), and every strategy $\hat{\mathbf{r}}_i$,

$$\hat{U}_i(\theta_0, F; \hat{\mathbf{r}}_i^{\mathrm{tr}}) \geq \hat{U}_i(\theta_0, F; \hat{\mathbf{r}}_i).$$

Proof of Proposition 3. Fix an agent i, a realized initial state (θ_0, F), and an arbitrary strategy $\hat{\mathbf{r}}_i$ in the outset mechanism. We construct a strategy $\tilde{\mathbf{r}}_i$ in the repeated-report environment that generates exactly the same allocation and transfer path whenever the other agents report truthfully. To this end, let $\hat{r}_{i0}(\theta_{i0}, F_i) = (\tilde{\theta}_{i0}, G_i)$. Define $\tilde{\mathbf{r}}_i$ as follows. In period 0, agent i sends the same message as in the outset mechanism: $\tilde{r}_{i0}(\theta_{i0}, F_i) \equiv \hat{r}_{i0}(\theta_{i0}, F_i)$. For each $t \geq 1$, let $\hat{h}_{it}$ be the reduced private history obtained from the repeated-report private history h_{it} by deleting, from every period s public report with $1 \leq s \leq t-1$, the kernel components that are repeated from period 0. We then set $\tilde{r}_{it}(h_{it}) \equiv \big(\hat{r}_{it}(\hat{h}_{it}), G_i\big)$. Thus $\tilde{\mathbf{r}}_i$ repeats in every period the same kernel report G_i chosen by $\hat{\mathbf{r}}_i$ in period 0, and otherwise mimics the payoff-type reports prescribed by $\hat{\mathbf{r}}_i$.

By construction of the outset mechanism, whenever agents $j \neq i$ report truthfully, the path of public histories generated under $\hat{\mathbf{r}}_i$ in the outset mechanism lifts period by period to the path of public histories generated under $\tilde{\mathbf{r}}_i$ in the repeated-report environment. Consequently, the allocations and transfers coincide pathwise across the two environments. It follows that $\hat{U}_i(\theta_0, F; \hat{\mathbf{r}}_i) = U_i(\omega_0; \tilde{\mathbf{r}}_i)$, where $\omega_0 = (\theta_0, F)$ and $U_i(\omega_0; \tilde{\mathbf{r}}_i)$ denotes agent i's expected discounted utility in the repeated-report mechanism when the other agents report truthfully.

By Lemma 1(i), $U_i(\omega_0; \tilde{\mathbf{r}}_i) \leq V_i^*(\omega_0, h_0)$. Moreover, the truthful strategy $\hat{\mathbf{r}}_i^{\mathrm{tr}}$ in the outset mechanism lifts, under the same construction, to the truthful strategy in the repeated-report environment, so that, by Lemma 1(ii),

$$\hat{U}_i(\theta_0, F; \hat{\mathbf{r}}_i^{\mathrm{tr}}) = U_i(\omega_0; \mathbf{r}_i^{\mathrm{tr}}) = V_i^*(\omega_0, h_0).$$

Combining the displayed relations yields $\hat{U}_i(\theta_0, F; \hat{\mathbf{r}}_i^{\mathrm{tr}}) \geq \hat{U}_i(\theta_0, F; \hat{\mathbf{r}}_i)$. Since i, (θ_0, F), and $\hat{\mathbf{r}}_i$ were arbitrary, the outset mechanism is ex-post incentive compatible at the outset. □

B Details from Section 4.3

Proof of Proposition 4. Fix a realized initial payoff-type profile θ_0. We first note that, given the agents' common observation of θ_0, truthful reporting remains a continuation equilibrium after every acquisition profile $F \in \mathcal{F}$, including profiles reached by unilateral deviations. Propositions 1 and 2, Proposition 3, and Corollary 1 establish incentive compatibility for every realized state (θ_0, F); in particular, the bounds they provide hold realization by realization (Lemma 1 bounds deviation payoffs pointwise in the realized types of the other agents). Consequently, an agent who observes θ_0 finds truthful reporting optimal after every acquisition profile. Hence, when studying the preliminary investment stage, it suffices to compute each agent's continuation payoff under truthful reporting after every acquisition profile, and to check deviations in the acquired kernel alone; a joint deviation $(\widehat{F}_i, \mathbf{r}_i)$ in kernel and reporting strategy is weakly dominated by the deviation $(\widehat{F}_i, \mathbf{r}_i^{\mathrm{tr}})$.

Under truthful reporting in the dynamic team mechanism, agent i's flow utility equals total flow surplus. Her net payoff from an acquisition profile F is therefore

$$\mathbb{E}_F \left[\sum_{t=0}^{\infty} \delta^t \sum_{j \in N} v_j\big(x_t^*(\theta_t, F), \theta_{jt}\big) \,\middle|\, \theta_0 \right] - c_i(F_i) = W(\theta_0, F) - c_i(F_i).$$

Under truthful reporting in the dynamic pivot mechanism, her net payoff is

$$M_i(\theta_0, F) - c_i(F_i) = W(\theta_0, F) - W_{-i}(\theta_{-i,0}, F_{-i}) - c_i(F_i).$$

Fix i, $F = (F_i, F_{-i}) \in \mathcal{F}$, and $\widehat{F}_i \in \mathcal{F}_i$. In the team mechanism, the change in agent i's net payoff from replacing F_i by $\widehat{F}_i$ is

$$W(\theta_0, \widehat{F}_i, F_{-i}) - c_i(\widehat{F}_i) - W(\theta_0, F_i, F_{-i}) + c_i(F_i) = \Phi(\theta_0, \widehat{F}_i, F_{-i}) - \Phi(\theta_0, F_i, F_{-i}),$$

where the costs of agents $j \neq i$ cancel. The same identity holds in the pivot mechanism because $W_{-i}(\theta_{-i,0}, F_{-i})$ is independent of F_i and therefore also cancels. Thus, under either mechanism, the acquisition game with truthful continuation play is an exact potential game with potential $\Phi(\theta_0, \cdot)$.

Since $F^* = F^*(\theta_0)$ maximizes $\Phi(\theta_0, \cdot)$, it constitutes a Nash equilibrium of the acquisition game followed by truthful reporting (Monderer and Shapley, 1996). As noted above, any joint deviation in acquisition and reporting is weakly dominated by one with the same acquisition choice and truthful reporting, and is therefore unprofitable. Since this holds for every realized θ_0, acquiring $F^*(\theta_0)$ and then reporting truthfully is an ex-post equilibrium at the outset. Along this equilibrium path, the mechanism implements the efficient allocation associated with $F^*(\theta_0)$, and $F^*(\theta_0)$ maximizes cost-adjusted social surplus by construction. $\square$

References

Acemoglu, D. and Lensman, T. (2024). Regulating transformative technologies. *American Economic Review: Insights*, 6(3):359–376.

Athey, S. and Segal, I. (2007). An efficient dynamic mechanism. Working paper, Stanford University.

Athey, S. and Segal, I. (2013). An efficient dynamic mechanism. *Econometrica*, 81(6):2463–2485.

Baron, D. P. and Besanko, D. (1984). Regulation and information in a continuing relationship. *Information Economics and Policy*, 1(3):267–302.

Bergemann, D. and Morris, S. (2005). Robust mechanism design. *Econometrica*, 73(6):1771–1813.

Bergemann, D. and Välimäki, J. (2002). Information acquisition and efficient mechanism design. *Econometrica*, 70(3):1007–1033.

Bergemann, D. and Välimäki, J. (2010). The dynamic pivot mechanism. *Econometrica*, 78(2):771–789.

Bergemann, D. and Välimäki, J. (2019). Dynamic mechanism design: An introduction. *Journal of Economic Literature*, 57(2):235–274.

Bergemann, D., van Ryzin, G., and Li, J. (2026). Supply chain coordination mechanism design: Consensus planning protocol meets Vickrey-Clarke-Groves mechanism. *arXiv preprint arXiv:2605.16695*.

Bertsekas, D. P. and Shreve, S. E. (1978). *Stochastic Optimal Control: The Discrete-Time Case*. Academic Press, New York.

Camara, M. K., Hartline, J. D., and Johnsen, A. (2020). Mechanisms for a no-regret agent: Beyond the common prior. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 259–270. IEEE.

Carroll, G. (2017). Robustness and separation in multidimensional screening. *Econometrica*, 85(2):453–488.

Carroll, G. (2019). Robustness in mechanism design and contracting. *Annual Review of Economics*, 11:139–166.

Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, Learning, and Games*. Cambridge University Press.

Chassang, S. (2013). Calibrated incentive contracts. *Econometrica*, 81(5):1935–1971.

Chung, K.-S. and Ely, J. C. (2007). Foundations of dominant-strategy mechanisms. *The Review of Economic Studies*, 74(2):447–476.

Csóka, E. (2025). Prior-free collusion-proof dynamic mechanisms. *arXiv preprint arXiv:2511.15727*.

Csóka, E., Liu, H., Rodivilov, A., and Teytelboym, A. (2023). A collusion-proof efficient dynamic mechanism. *Available at SSRN 4419623*.

Han, Q., Simchi-Levi, D., Tan, R., and Zhao, Z. (2025). Multi-agent adaptive mechanism design. *arXiv preprint arXiv:2512.21794*.

Koçyiğit, Ç., Iyengar, G., Kuhn, D., and Wiesemann, W. (2020). Distributionally robust mechanism design. *Management Science*, 66(1):159–189.

Koh, A. and Sanguanmoo, S. (2025). Robust technology regulation. *MIT, Job Market Paper*.

Laffont, J.-J. and Tirole, J. (1993). *A Theory of Incentives in Procurement and Regulation*. MIT Press.

Libgober, J. and Mu, X. (2021). Informational robustness in intertemporal pricing. *The Review of Economic Studies*, 88(3):1224–1252.

Monderer, D. and Shapley, L. S. (1996). Potential games. *Games and Economic Behavior*, 14(1):124–143.

Pavan, A., Segal, I., and Toikka, J. (2014). Dynamic mechanism design: A myersonian approach. *Econometrica*, 82(2):601–653.

Penta, A. (2015). Robust dynamic implementation. *Journal of Economic Theory*, 160:280–316.

Quint, D. and Kojima, F. (2026). Symmetric equilibrium in pre-auction investment. *Available at SSRN*.

Rogerson, W. P. (1992). Contractual solutions to the hold-up problem. *The Review of Economic Studies*, 59(4):777–793.

Wilson, R. (1987). Game-theoretic analyses of trading processes. In Bewley, T. F., editor, *Advances in Economic Theory: Fifth World Congress*, pages 33–70. Cambridge University Press.