EFFICIENT DYNAMIC MECHANISM DESIGN WITHOUT A COMMON PRIOR

By

Dirk Bergemann and Marek Bojko

July 2026

COWLES FOUNDATION DISCUSSION PAPER NO. 2540

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[‡]

July 3, 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 stochastic kernel governing the evolution of her own type and may hold arbitrary beliefs about the kernels of others. We extend the agents' type space to include the kernel itself 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 can be elicited only once, at the outset, and that the same mechanisms induce the efficient private acquisition of the stochastic 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. 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. The classic approach to such problems assumes that all parties share a common prior over payoff-relevant states and their stochastic evolution. While analytically convenient, this assumption rarely holds in practice. A retailer's forecast of demand is built from proprietary point-of-sale, search, and customer-relationship data. A vendor's beliefs about component availability rest on its own relationships with upstream suppliers and on its assessment of geopolitical risk. There is little reason to expect these diverse sources of information to be derived from any common probabilistic model. The same observation applies more broadly: in dynamic procurement, dynamic regulation (Laffont and Tirole, 1993; Acemoglu and Lensman, 2024; Koh and Sanguanmoo, 2025), and dynamic public-good provision, the stochastic models governing payoffs are typically the private property of the agents who hold them.

This paper clarifies the role of the common prior assumption in efficient dynamic mechanism design. We work in an environment with independent private values in which each agent privately knows the stochastic kernel that governs the evolution of her own type. We do not require the other agents, or the designer, to know that kernel; nor do we impose a common prior over kernel profiles. Higher-order beliefs need not be specified. We consider mechanisms that elicit both the stochastic kernels and the agents' realized valuations, treating the kernel as part of the agent's type. Under the interpretation that the efficient allocation is the policy maximizing expected surplus in the stochastic environment specified by the agents' individual priors, rather than a contingent plan derived from a fixed common prior, the analysis also accommodates subjective beliefs. In this respect, the framework remains applicable even when agents' priors are subjective or evolve over time.[1]

We first show that the dynamic team mechanism of Athey and Segal (2013) and the dynamic pivot mechanism of Bergemann and Välimäki (2010) continue to implement the efficient allocation when agents' transition kernels are elicited as part of their private information (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 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 explain in Section 4.1, the results also follow directly from the implementation results of Athey and Segal (2013) and Bergemann and Välimäki (2010) applied to suitably constructed type spaces and transition laws. Explicitly distinguishing payoff types from stochastic kernels makes transparent that efficient dynamic implementation does not require common knowledge of the stochastic environment governing the evolution of agents' payoff types.

We then use this observation in two ways. First, Proposition 3 and Corollary 1 show that, when

[1] We discuss this point further in Section 4.4.

kernels are persistent, they do not have to be reported repeatedly in every period. The mechanism can elicit them 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.[2] This matters because the stochastic laws used by agents are often not simply given. They may be the product of costly investments in experimentation or forecasting.

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 robust mechanism design literature following Wilson (1987), Bergemann and Morris (2005), Chung and Ely (2007), Carroll (2017, 2019), and Koçyiğit et al. (2020); we consider dynamic environments. Chassang (2013), Penta (2015), Libgober and Mu (2021), Koh and Sanguanmoo (2025), and Csóka (2025), among others, consider robustness in dynamic mechanism design; we focus on a distinct and complementary question of eliciting priors.

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. In Section 2, we present the formal model. In Section 3, we provide our main results for the dynamic team and dynamic pivot mechanisms. In Section 4, we consider alternative timing assumptions and interpretations of the model and results. 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 $x_t \in X$, for a compact metric space X, the current monetary transfer $p_{it} \in \mathbb{R}$, and the private *preference type* $\theta_{it} \in \Theta_i$, for a measurable set Θ_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

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

[2] Athey and Segal (2007), an earlier working paper version of Athey and Segal (2013), showed that the dynamic team mechanism induces efficient investment.

Agents share a discount factor $\delta \in (0,1)$. To ensure that all expected payoffs below are well defined, we assume that there exists $K > 0$ such that $\|v_i\|_\infty < K$ for every i; standard measurability assumptions on v_i are imposed throughout.

The preference 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 $F \in \mathcal{F}$. Conditional on the current state and allocation, agents' type transitions are independent across agents, so the aggregate transition kernel is

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

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 preference 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: the modeler does not specify 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. We assume throughout that the Bellman operator admits a measurable optimal policy.[3] Stated recursively, W solves the Bellman equation

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

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. 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. The private history of agent i in period t further records the agent's realized types,

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

with set of private histories 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.

A pure reporting strategy of agent i is a 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

[3] Along with the assumptions stated above, it is sufficient to assume v_i is continuous, Θ is a Polish space, $\mathcal{F}$ is finite, and F_i is Feller in (θ_i, x) for each agent i and $F_i \in \mathcal{F}_i$. In this case, the Bellman operator maps bounded continuous functions into bounded continuous functions and is a contraction under the sup norm. Hence its fixed point W is bounded and continuous. Existence of a measurable maximizer then follows by the measurable maximum theorem.

strategies $\mathbf{r}_{-i}$ and types are drawn from the true kernel profile F, is

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

We can write the agent's continuation value recursively as

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

The value function $V_i^{\mathbf{r}_{-i}}$ 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 follow strategies $\mathbf{r}_{-i}$. Denote the value function when the rest of the agents report truthfully by V_i^*. Our equilibrium concept is periodic ex-post equilibrium.

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

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

When $(\mathbf{x}^*,\mathbf{p})$ is periodic ex-post incentive compatible, no agent has an incentive to deviate from truth-telling after learning the contemporaneous types of the other agents. Note that these contemporaneous types of the other agents contain information about the future evolution of their payoff types. 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 from the corresponding implementation results in 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, history h_t, and every profile $\omega_t = (\theta_{it}, F_i)_{i\in N}$

$$p_i(\omega_t) = -\sum_{j\neq i} v_j\big(x_t^*(\omega_t), \theta_{jt}\big),$$

where we suppress the public history in the transfer. In words, each agent is paid in each period the gross flow surplus that the other agents receive at the efficient allocation.

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

Proof. By the one-stage deviation principle, it suffices to consider a deviation by agent i in a single period t, reporting $r_{it} \neq \omega_{it}$ while reporting truthfully thereafter. Under the team transfer, the agent's continuation payoff from such a deviation is

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

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. $\square$

A report of F_i affects agent i's payoff only through the induced allocation path. Continuation expectations are evaluated using the true kernel.

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, as derived by Bergemann and Välimäki (2010), the flow marginal contribution is

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

where $x^*_{-i,t}$ denotes the socially efficient allocation in the economy without agent i. The pivot transfer is then defined by

$$p_i(\omega_t) = v_i(x^*_t(\omega_t), \theta_{it}) - m_i(\omega_t).$$

Thus agent i's flow utility under the mechanism 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. By the principle of optimality, 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:

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

for all $r_{it} \in \Omega_i$ and $\omega_t \in \Omega$. By construction of p_i, the left-hand side represents the marginal contribution of agent i. Plugging in for M_i on the right-hand side, the inequality above is equivalent to

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

Further, note that

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

In particular, the latter summand does not depend directly on the report of F_i, except through the induced current allocation x^*_t. Plugging in for $p_i(r_{it}, \omega_{-i,t})$ and collecting terms, we obtain

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

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

The terms corresponding to the economy without agent i do not depend directly on agent i's reported kernel, except through the current allocation. They may, however, depend on the kernels reported by the other agents. In the desired equilibrium, those kernels are reported truthfully. Hence the beliefs used by agent i coincide with the beliefs embedded in the transfers.

4 Discussion and Extensions

4.1 Common Prior on Expanded Type Spaces

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 determines how the agent's payoff type evolves in the future. In this section, we clarify to what extent our environment can be represented as a standard dynamic private-value environment with a common prior on the expanded type space. For simplicity of exposition, we assume $\mathcal{F}$ is finite.

On this expanded space, for $\omega_i = (\theta_i, F_i), \omega_i' = (\theta_i', F_i') \in \Omega_i$ and $x \in X$, define agent i's transition kernel by

$$\mu_i(d\omega_i' \mid \omega_i, x) = F_i(d\theta_i' \mid \theta_i, x) \cdot \mathbb{1}_{F_i' = F_i}.$$

Thus the payoff component evolves according to the kernel contained in the current type, while the kernel component is persistent. The aggregate transition kernel μ is given by a product of individual transition kernels. This is a commonly known kernel on the expanded type space Ω. We have thus "markovized" the environment (Athey and Segal, 2013).

The incentive arguments of Athey and Segal (2013) and Bergemann and Välimäki (2010) can be applied directly to the expanded type space. The dynamic team mechanism of Athey and Segal (2013) makes each agent a residual claimant of total surplus on the expanded state space. Similarly, the dynamic pivot mechanism of Bergemann and Välimäki (2010) gives each agent her dynamic marginal contribution on the expanded state space. Propositions 1 and 2 can therefore be viewed as applications of these results after replacing the type θ_i by the expanded type (θ_i, F_i).

Taken together, the preceding section and the discussion above clarify the role of the common prior assumption. We do not require the designer to know the realized kernel profile F, nor do we require a common prior over possible kernel profiles. In particular, there need not be a prior over the initial type profile ω_0, nor an exogenous initial decision. For the dynamic team and dynamic pivot mechanisms, a commonly known stochastic kernel for types is without loss once the type space is expanded sufficiently. Nevertheless, explicitly distinguishing payoff types from stochastic kernels within the agents' types ω remains conceptually useful, as it makes transparent why stochastic kernels can be elicited from the agents.

The distinction is also useful for the results below: persistence of the kernel component is what permits one-time elicitation at the outset, and the interpretation of F_i as a stochastic model is what makes it meaningful to study costly investment.

4.2 Eliciting Priors 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 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 a 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, given truthful reporting of other agents, every realized initial state (θ_0, F), and every reporting strategy $\mathbf{r}_i$, reporting truthfully yields a weakly higher expected payoff than following $\mathbf{r}_i$.

Proposition 3 (One-time elicitation of kernels)**.** *Suppose that the 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.*

The proof sketch is as follows. 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 priors 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. We analyze the investment game pointwise in the realized initial payoff-type profile θ_0. Let $c_i : \mathcal{F}_i \to \mathbb{R}_+$ denote the cost of acquiring kernel $F_i \in \mathcal{F}_i$; the cost is paid before the dynamic mechanism begins, after which F_i becomes agent i's private kernel as in the baseline model. Fix the realized initial payoff-type profile θ_0 and define

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

The efficient acquisition problem is to maximize Φ over $F \in \mathcal{F}$; we assume that a solution F^* exists. 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 $(F_i, \mathbf{r}_i)$ of an alternative kernel and reporting strategy, acquiring F_i^* and reporting truthfully yields a weakly higher expected payoff than $(F_i, \mathbf{r}_i)$, given that other agents acquire F_{-i}^* and report truthfully.

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.

The key observation is that 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 fixed F_{-i}, agent i's truthful expected payoff gain from choosing $\widehat{F}_i$ rather than F_i is

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

where Φ denotes cost-adjusted social surplus. This mirrors the observation of Quint and Kojima (2026) in a static environment. It follows that the acquisition game induced by either mechanism is a potential game with potential function Φ. Consequently, every global maximizer of Φ constitutes a Nash equilibrium of the acquisition stage, provided agents report truthfully in the continuation game.

4.4 Evolving Subjective Priors

Rather than viewing F_i as a fixed primitive of the environment, we may also interpret the kernel reported in a given period as agent i's current assessment of how her future payoff types will evolve. Under this interpretation, the same kernel need not persist. An agent may learn about new contingencies or revise her forecasting model. The mechanism simply conditions the continuation problem on the kernel profile that is currently reported.

With this interpretation, the welfare benchmark also changes. Efficiency is no longer defined relative to a single kernel profile reported at period 0, but relative to the kernel profile currently in force. In each period, the planner chooses the allocation that maximizes current surplus plus discounted continuation surplus under the current kernel profile; if that profile changes in a later period, the planner resolves the dynamic problem using the new kernels. The socially efficient allocation is therefore best understood as a sequence of conditionally efficient allocations rather than as a single plan derived from an immutable prior. In this sense, our analysis remains useful even when priors evolve over time: it identifies the allocation that is efficient for the kernels currently held by the agents, without requiring a Bayesian model of how those kernels are updated.

5 Concluding Remarks

This paper extends the Wilson doctrine (Wilson, 1987; Bergemann and Morris, 2005) to dynamic mechanism design. We show that, under independent private values, efficient implementation requires neither common knowledge of agents' beliefs nor common knowledge of the stochastic environment. More broadly, the results identify a class of environments in which efficient dynamic implementation is robust to uncertainty about both agents' beliefs and the underlying stochastic environment. Relaxing common prior assumptions in this way is important for several applications, such as 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).

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*.

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.

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.

A Details from Section 4.2

Fix a direct mechanism $(\mathbf{x}^*, \mathbf{p})$ in the environment of Section 2. 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

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

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.

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 $\hat{\mathbf{r}}_i^{\text{tr}}$ denote the truthful strategy, defined by

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

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 $\hat{\mathbf{r}}_i$,

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

Finally, we prove Proposition 3.

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 allocation and transfer at every period 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.

Now consider the repeated-report environment. By definition of V_i^*, for every public history h_t and current type ω_t, the quantity $V_i^*(\omega_t, h_t)$ is the maximal continuation payoff available to agent i when the other agents report truthfully. Periodic ex-post incentive compatibility implies that, at every such (ω_t, h_t), the truthful report ω_{it} attains the maximand that defines $V_i^*(\omega_t, h_t)$. Therefore truthful continuation yields the value $V_i^*(\omega_t, h_t)$ from every history and, in particular, a weakly higher payoff than every alternative strategy from period 0 onward. Hence, $U_i(\omega_0; \tilde{\mathbf{r}}_i) \leq V_i^*(\omega_0)$. Finally, the truthful strategy in the outset mechanism corresponds exactly to the truthful strategy under the same lifting argument. Hence, $V_i^*(\omega_0) = \hat{U}_i(\theta_0, F; \hat{\mathbf{r}}_i^{\mathrm{tr}})$. Combining the displayed equalities and inequalities yields the desired inequality. Since i, (θ_0, F), and $\hat{\mathbf{r}}_i$ were arbitrary, the outset mechanism is ex-post incentive compatible at the outset. $\square$

B Details from Section 4.3

Proof of Proposition 4. By Propositions 1, 2, 3, and Corollary 1, truthful reporting is optimal after every acquired kernel profile F. Hence, when studying the preliminary investment stage, it suffices to compute each agent's continuation payoff under truthful reporting after acquisition.

First consider the dynamic team mechanism. Agent i's flow utility under truthful reporting is

$$\sum_{j \in N} v_j(x_t^*(\theta_t, F), \theta_{jt}).$$

Therefore agent i's gross continuation payoff, before subtracting the acquisition cost, is

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

Thus agent i's net payoff from acquiring F_i, when the other agents acquire F_{-i}, is

$$W(\theta_0, F) - c_i(F_i).$$

Now consider the dynamic pivot mechanism. Let

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

denote agent i's dynamic marginal contribution. By construction of the dynamic pivot transfers, agent i's flow utility is the flow marginal contribution m_i, and hence her gross continuation payoff under truthful reporting is

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

Therefore her net payoff from acquiring F_i, when the other agents acquire F_{-i}, is

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

We now verify the investment incentive. Fix the efficient investment profile F^*. For any alternative acquisition choice $\widehat{F}_i \in \mathcal{F}_i$, optimality of F^* gives

$$W(\theta_0, F^*) - \sum_{j \in N} c_j(F_j^*) \geq W(\theta_0, \widehat{F}_i, F_{-i}^*) - c_i(\widehat{F}_i) - \sum_{j \neq i} c_j(F_j^*).$$

Canceling the costs of agents $j \neq i$, we obtain

$$W(\theta_0, F^*) - c_i(F_i^*) \geq W(\theta_0, \widehat{F}_i, F_{-i}^*) - c_i(\widehat{F}_i).$$

This is exactly the no-profitable-deviation condition for agent i in the preliminary investment stage under the team mechanism. For the pivot mechanism, the same inequality applies because $W_{-i}(\theta_{-i,0}, F_{-i}^*)$ is independent of F_i.

Finally, after any acquisition profile, truthful reporting is a continuation equilibrium of the dynamic team mechanism and of the dynamic pivot mechanism. Thus an agent cannot gain by combining an investment deviation with a subsequent reporting deviation. Since i and $\widehat{F}_i$ were arbitrary, no agent has a profitable deviation from F^* in the acquisition stage. Therefore the strategy profile in which agents acquire F^* and then report truthfully is an ex-post equilibrium. Along this equilibrium path, the mechanism implements the efficient allocation for the acquired kernel profile F^*, and F^* maximizes cost-adjusted social surplus by construction. □