A Collusion-Proof Efficient Dynamic Mechanism*

Endre Csóka[†], Heng Liu[‡], Alexander Rodivilov[§], Alexander Teytelboym[¶]

May 14, 2024

Abstract

We characterize an efficient, budget-balanced, individually rational, and collusion-proof mechanism in a dynamic environment with private values and independent types. Agents observe private information, and public decisions are made over time. Each agent guarantees himself an expected utility level by being truthful, and those guaranteed utility levels add up to the highest total ex-ante expected surplus. As a result, all equilibria sustained by our mechanism are efficient and utility-equivalent and the mechanism is robust to any reasonable collusive agreement. Our results remain intact if agents choose private actions and observe true past types of other agents. The properties of our mechanism stand in stark contrast to the Dynamic Pivot Mechanism (Bergemann and Välimäki, 2010) and Balanced Team Mechanism (Athey and Segal, 2013), which, as we show, might admit only inefficient equilibria after the elimination of weakly dominated strategies. We also construct a modified mechanism that approximately achieves the same property in environments without transfers.

Keywords: Dynamic mechanism design, dynamic incentive compatibility, collusion-proofness, budget balance.

JEL classification: D21, D71, D81, D82, D86.

*This paper merges and supersedes two papers: (i) Csóka (2015), except his model and results in a setting with adverse selection, and (ii) Liu (2022). We are thankful for the helpful comments of Eric Bahel, Paulo Barelli, Tilman Börgers, Jon Eguia, Hari Govindan, Johannes Hörner, David Miller, Debasis Mishra, Larry Samuelson, Ilya Segal, Ludvig Sinander, Satoru Takahashi, and participants at the 2022 Conference on Economic Design (University of Padova), the 2022 International Conference "Game Theory and Applications", 2022 Stony Brook International Conference on Game Theory, and 2023 Canadian Economic Theory Conference.

[†]Alfréd Rényi Institute of Mathematics, Budapest, Hungary. *Email*: csoka.endre@renyi.hu.

[‡]Department of Economics, Rensselaer Polytechnic Institute, USA *Email*: hengliu29@gmail.com.

[§]School of Business, Stevens Institute of Technology, USA. *Email*: arodivil@stevens.edu.

[¶]Department of Economics, University of Oxford, UK. *Email*: alexander.teytelboym@economics.ox.ac.uk.

1 Introduction

Dynamic mechanism design addresses questions of information elicitation and resource allocation in environments where agents' private information evolves stochastically over time and can be affected by allocation decisions (Bergemann and Välimäki, 2019). The framework captures many important economic settings, such as dynamic public good provision and nonlinear pricing.

In this paper, we are interested in dynamic mechanisms that achieve a socially efficient allocation. Our main contribution is a new dynamic mechanism that is efficient, individually rational, budget-balanced, and, crucially, collusion-proof. The main property of our mechanism is that every agent can obtain a guaranteed expected utility level even if all other agents conspire to minimize his utility—these guaranteed utility levels exactly add up to the expected efficient surplus. As a result, not only is *every* Bayes-Nash equilibrium of our mechanism efficient, but it is resistant to any conceivable type of collusion.

The properties of our mechanism stand in stark contrast to two classic efficient mechanisms that sustain truth-telling in a Perfect Bayesian Equilibrium: the Dynamic Pivot Mechanism (Bergemann and Välimäki, 2010) (which is, however, not budget-balanced) and the Balanced Team Mechanism (Athey and Segal, 2013) (which, however, only satisfies weaker participation constraints). It is not difficult to see that since these mechanisms are dynamic versions of the Vickrey-Clarke-Groves (VCG) and the Arrow/d'Aspremont–Gérard–Varet (AGV) mechanisms respectively, they might be susceptible to collusion. However, we show that it is possible for none of the truth-telling Perfect Bayesian Equilibria of either mechanism to survive iterated elimination of weakly dominated strategies. Our finding suggests that efficient equilibria in these mechanisms are surprisingly fragile, and our mechanism offers a more robust alternative without sacrificing any of the desirable properties of the Dynamic Pivot and Balanced Team Mechanisms.

We consider a general dynamic environment à la Athey and Segal (2013). Every period, agents observe their private types and a public state; after types are reported a public decision is made. The type distribution of a private type can be affected by public decisions. Otherwise, the agents' types evolve independently. Agents' utility functions are quasilinear and can depend on private signals and public decisions arbitrarily.

Every period, each agent makes a public report of his private type. A mechanism consists of a decision policy that determines a public decision and a payment rule that determines transfers to the agents, both as a function of the reported types. The mechanism is efficient if the decision policy maximizes the expected sum of all agents' payoffs. The mechanism is budget-balanced if, in every period, the sum of all transfers to all agents is zero.

We now describe our *Guaranteed Utility Mechanism* (**GUM**). The high-level idea is the following. When multiple reports arrive, they jointly impose an externality on each agent because they change the public decision. As usual, externality refers to the change in the total "anticipated" (i.e., the expectation is calculated from believing the past reports and assuming truthfulness in the future) payoffs of others. The key decision is how to calculate the share of an individual report in this joint externality. In our mechanism, these shares are *additive*. Namely, if multiple reports are made, the joint externality imposed by simultaneously updating these reports equals the sum of the calculated individual externalities. And each agent i pays the transfer to j equal to this externality j imposed on i. If the externality is additive, then members of a collusive cartel cannot exploit the difference between the sum of their calculated individual externalities and the true joint effect of simultaneous updating of their reports.

The GUM has two key features that distinguish it from other mechanisms. First, when calculating the externality induced by an agent's report, we process all the reports *sequentially* according to the ordering $i = 1, ..., N$.[1] For an example with three agents, when we process the report of agent 2 in round t, we calculate the anticipated total payoffs of the agents given the round-t reported type of agent 1, and the round-$(t-1)$ reported types of agents 2 and 3. Then we calculate how agents' payoffs change by the update to the round-t report of agent 2.

The second difference is that the externality payments are *bilateral*: every agent j receives the transfer from i which compensates j by the externality imposed on j by i. In our example, assume that the update increases the anticipated total payoff of agent 1 by 20 and decreases the anticipated total payoff of agent 3 by 5. Then the transfers are $(-20, 15, 5)$, namely, agent 1 pays 20 to agent 2 and agent 2 pays 5 to agent 3. By construction, this ensures that GUM is budget-balanced.

In a truth-telling equilibrium of the GUM, every agent obtains his maxmin utility and agents' maxmin utilities sum up to the maximum ex-ante surplus, so the mechanism is efficient. As a consequence, all Bayes-Nash equilibria of GUM are efficient and utility-equivalent (Proposition 1). Although truth-telling is not a dominant strategy[2], there is essentially nothing for the agents to gain from misreporting their type under the GUM. Indeed, we show that the GUM is *collusion-proof* in a very strong sense, i.e., whenever the colluding agents

[1] In Section H, we show that the mechanism can be symmetrized so that it is independent of the ordering; payments are then equal to the Shapley values where the externality plays the role of the value of the coalition.

[2] Efficiency is incompatible with dominant-strategy implementation in dynamic settings (Bergemann and Välimäki, 2019, Section 3.3).

can jointly coordinate their reports and exchange transfers within the coalition (Theorem 1). Intuitively, GUM is collusion-proof because every honest agent obtains an expected utility guarantee regardless of the strategies of other agents, and those guarantees add up to the socially efficient surplus.

Having established the efficiency and collusion-proofness of the GUM, we turn to participation constraints and other properties. First, we show that GUM can be enhanced to satisfy a dynamic version of an interim participation constraint, explored by Athey and Segal (2013), which allows agents to exit every round having learned their type and re-enter the mechanism later (Section 6.1). While we used a standard dynamic mechanism design setup, our results can be extended to a more general environment. For example, our results remain intact if agents observe informative signals of the true past types of other agents (Section 6.2), and, more interestingly, even if agents can choose private actions (Section 6.3).[3]

The payment rule in GUM differs from the payment rules for two classic efficient dynamic mechanisms: the Dynamic Pivot Mechanism (DPM) due to Bergemann and Välimäki (2010) and the Balanced Team Mechanism (BTM) due to Athey and Segal (2013).

The DPM is the dynamic version of the VCG mechanism. In DPM, each agent gets a reward equal to his flow marginal contribution. That is, the expected transfer from every agent is equal to the dynamic externality he imposes on others. These pivot-like payments satisfy ex-post participation constraints and period ex-post incentive constraints (i.e., robustness to observing same-period types which the GUM does not satisfy) are satisfied for all histories, but are not budget-balanced and do not allow for private actions.

The BTM is the dynamic version of the AGV mechanism. Like the GUM, it is efficient, budget-balanced, and even allows for private actions. However, there are two main differences between the GUM and the BTM. First, the way the incentive term is calculated. In the BTM, the externality on the agent j's payoff is calculated by updating *only* the report of agent i at period t while keeping the reports of all the other agents at the period-$(t-1)$ value. In our mechanism, all the agents' reports are updated sequentially at period t according to the order $i = 1, ..., N$. Second, the way the transfers are paid. In the BTM, each agent i's incentive term is paid *equally by all the other agents* $j \neq i$. In the GUM, every agent j is compensated for the externality imposed on him by i's updated report. As our payment rule makes each agent accountable for the externality his report imposes on other agents, the mechanism remains budget-balanced without distorting the truth-telling incentives.

It is not surprising that neither the DPM nor the BTM is collusion-proof. Crucially, however, both in the DPM and the BTM, the externality is not additive. As a result, we

[3]In Appendix H, we also discuss how to make GUM symmetric and the connection to Shapley values.

	DPM	BTM	GUM
Efficiency and truth-telling			
existence of an efficient PBE	YES	YES	YES
existence of an efficient PBE that survives iterated elimination of weakly dominated strategies	NO (Sec. 5.2)	NO (Sec. 5.1)	YES
all PBEs are efficient	NO	NO	YES
all Bayesian Nash equilibria are efficient	NO	NO	YES (Prop. 1)
Collusion			
collusion-proof (Sec. 4.2)	NO	NO	YES (Thm. 1)
Properties, robustness, extensions			
balanced budget	NO	YES	YES
exiting and re-entering are allowed (Sec. 6.1)	YES	NO	YES
robustness to observing past true types (Sec. 6.2)	YES	NO	YES
robustness to observing same-period types	YES	NO	NO
private actions (Sec. 6.3)	NO	YES	YES

Table 1: The properties of the Dynamic Pivot Mechanism (DPM, due to Bergemann and Välimäki (2010), the Balanced Team Mechanism (BTM, due to Athey and Segal (2013), and the Guaranteed Utility Mechanisms (GUM, this paper). PBE = Perfect Bayesian Equilibrium.

show that under both DPM and BTM, it is possible that none of the truth-telling equilibria survive iterated elimination of weakly dominated strategies and therefore the only surviving equilibria in these mechanisms are inefficient (Section 5). Such fragility sounds caution for any designer who is interested in applying these mechanisms in practice. Table 1 summarizes the differences between the DPM, the BTM, and the GUM.

Our mechanism points out a subtle link between dynamic collusion-proof mechanisms and the approximately efficient Bayesian mechanisms with multiple allocation decisions without transfers considered by Jackson and Sonnenschein (2007).[4] Specifically, in Section 7, we describe a transfer-free modification of GUM that achieves the same results as Jackson and Sonnenschein (2007) but with tighter error bounds. In particular, the transfer-free modification of GUM in which a truthful agent secures a payoff that converges to the ex-ante target level under the efficient decision policy. The basic idea is as follows. We modify GUM so that each agent is given a "budget." In every period, we calculate the sum of externalities imposed by an agent from the beginning up to that period. We then "punish"

[4] Also see Escobar and Toikka (2013), Renou and Tomala (2015), and Ball et al. (2022).

the agent if the sum exceeds his budget by replacing all his future reports with randomly generated ones. Our implementation is equivalent to the linking mechanism if agents face tight budget constraints but with higher error bounds.

The rest of the paper is organized as follows. In Section 2, we briefly discuss further related literature. In Section 3, we describe the setup. The Guaranteed Utility Mechanism and its key properties are in Section 4. In Section 5, we discuss the additive externality property of the GUM and point out key weaknesses of the DPM and BTM. The additional properties and extensions of the GUM are discussed in Section 6. In Section 7, we present a modification of our base model without transfers. Section 8 is a conclusion.

2 Related Literature

In an independent paper, Safronov (2018) introduced the special case of our mechanism in a static setting.[5] He argued that the mechanism is "coalition-proof." In his definition, colluding agents can coordinate reports and transfer money within the coalition. The types of agents within the coalition are commonly known. In Appendix C, we show that Safronov's notion of "coalition-proofness" is strictly weaker than our definition of collusion proofness even in a static setting and is susceptible to uncertainty about formed collusive coalitions. Although Safronov's mechanism coincides with GUM in a static setting, maintaining the attractive properties of GUM in a dynamic setting and achieving subgame perfection requires a far more complex construction. The reason is that, in dynamic settings, the distribution of future types used to calculate externalities evolves as a function of both agents' reports and public decisions. This distribution based on the reported types might differ from the true (i.e., based on true types) distribution if some agents are not truthful in the past. This feature is absent in a static environment. We overcome this challenge by exploiting a key martingale property of the anticipated continuation payoffs. Specifically, if an agent is truthful, for any period, the sum of his anticipated payoff and all the transfers up to that period is a martingale. As reports are updated sequentially within each period, our mechanism can be decomposed into a sequence of bilateral interactions. Our construction extends this insight to dynamic environments with general transition dynamics. In particular, we illustrate that it is insufficient for the agents to pay for the changes in the flow of expected utilities in each period due to the intertemporal correlation of types. Thus, the guaranteed utility mechanism requires that in each period, each agent pays the change in the continuation utilities of others conditional on the reports in the previous period. As the proof shows, it calibrates the flow

[5] Antecedents of these ideas include Crémer and Riordan (1985).

utility of an agent in every period to his ex-ante expected flow utility as long as he is always truthful.

A large literature, recently surveyed by Bergemann and Välimäki (2019), studies efficient mechanisms in a dynamic environment. Earlier work assumed that the agents did not learn any information relevant for the future after a mechanism they participate in has started (see, e.g., Atkeson and Lucas (1992) and Fudenberg et al. (1994), Wang (1995) and Miller (2011)).[6] thereby avoiding explicitly addressing the problems of contingent deviations in a dynamic setup. This assumption was then relaxed in several papers (Athey and Bagwell, 2008; Escobar and Toikka, 2013; Pavan et al., 2014; Golosov et al., 2014; Battaglini and Lamba, 2019).

Our results relate to the repeated implementation literature that studies full implementation by general (i.e., non-direct) mechanisms in environments with changing preferences. Lee and Sabourian (2011) and Mezzetti and Renou (2017) examine full implementation under complete information, and Lee and Sabourian (2013) considers incomplete information settings. Our collusion-proofness criterion can be viewed as full implementation in equilibrium utilities. The main difference is that we only consider direct mechanisms with monetary transfers in an independent private value environment with general transitions of preferences.

Finally, this paper is related to the literature on collusion-proof static mechanisms but with a different focus. Starting from Laffont and Martimort (1997) and Laffont and Martimort (2000), and the subsequent important contributions by Che and Kim (2006) and Che and Kim (2009), this literature studies whether and how the agents' ability to sign collusive agreements, usually modeled as side contracts among the agents, can undermine the principal's objective, such as revenue maximization, in static mechanism design problems. In contrast, this paper studies collusion-proofness in dynamic settings, where agents may collude by employing complex history-dependent strategies.

3 The Setup

There is a set $\mathcal{I} = \{1, ..., N\}$ of agents. Time is discrete, and the number of periods is finite.[7] The time periods are indexed by $t \in \{0, 1, ..., T\}$. In each period t, all the agents observe the

[6] One can avoid the challenge of contingent deviations by either considering a single agent with private information (e.g., Courty and Hao (2000), Battaglini (2005), Khalil et al. (2020), Rodivilov (2022), and Khalil et al. (2024)), or by allowing only an information structure that is independent across periods.

[7] We use a setup similar to the one in Athey and Segal (2013). The main difference is that Athey and Segal (2013) consider an infinite (countable) number of periods, whereas the number of periods is finite in our model. In addition, we assume there is no discounting. All of our results remain intact in the infinite-horizon environment.

(verifiable) public state $\theta_t^0 \in \boldsymbol{\Theta^0}$, and each agent $i \in \mathcal{I}$ privately observes his type $\theta_t^i \in \boldsymbol{\Theta^i}$. We denote the state space by $\boldsymbol{\Theta} = \prod_{i=0}^{N} \boldsymbol{\Theta^i}$. The initial state $\theta_0 \in \boldsymbol{\Theta}$ is fixed. After the state $\theta_t \in \boldsymbol{\Theta}$ is realized, a public decision $x_t \in \boldsymbol{X}$ is made, and each agent i collects a transfer $y_t^i \in \mathcal{R}$. All the described sets are finite.[8]

We now describe the transition probabilities and the distribution of subsequent states. The transition probability function is denoted by $\mu : \boldsymbol{\Theta} \times \boldsymbol{X} \to \Delta(\boldsymbol{\Theta})$. That is, given any state $\theta_t \in \boldsymbol{\Theta}$ and a public decision $x_t \in \boldsymbol{X}$ at period t, the $(t+1)$-period state is a random variable distributed according to the probability measure $\mu(\theta_t, x_t) \in \Delta(\boldsymbol{\Theta})$.[9] In addition, we assume that *types are independent*. That is, the transition probability measure μ can be written as

$$\mu(\theta, x) = \mu^0(\theta^0, x) \times \prod_{i=1}^{N} \mu^i(\theta^0, \theta^i, x), \tag{1}$$

where $\mu^0 : \boldsymbol{\Theta^0} \times \boldsymbol{X} \to \Delta(\boldsymbol{\Theta^0})$ and $\mu^i : \boldsymbol{\Theta^0} \times \boldsymbol{\Theta^i} \times \boldsymbol{X} \to \Delta(\boldsymbol{\Theta^i})$.[10] Intuitively, it implies that conditional on public information, an agent's private type does not affect the distribution of the current and future types of other agents.[11] The independent types assumption also implies that the agents cannot observe the past types of other agents. In Section 6.2, we consider a relaxation of this assumption and allow the agents to observe the true past types.

Utilities. We denote by $y_t^i \in \mathcal{R}$ a transfer to agent $i \in \mathcal{I}$ at period t. For a given sequences of types $\theta^i = \{\theta_t^i\}_{t=0}^T$, public decisions $x = \{x_t\}_{t=0}^T$, and transfers $y^i = \{y_t^i\}_{t=1}^T$, the utility of each agent $i \in \mathcal{I}$ is given by

$$U^i(\theta^i, x, y^i) = \sum_{t=0}^{T} \left[u^i(\theta_t^i, x_t) + y_t^i \right]. \tag{2}$$

That is, we have an environment with *private values* since the per-period payoff function $u^i(\theta_t^i, x_t)$ of each agent directly depends only on his private type and public decisions. Intuitively, this means that each agent can calculate his utility by only observing the public decision and his type, even if he is uncertain about other agents' types. Note that although the agents' utilities are quasilinear, we assume they can depend on private signals and public

[8]Our results can be extended to a setup with compact measurable spaces and upper semi-continuous payoff functions.

[9]See Bergemann and Välimäki (2019) for a discussion of Markovian transition functions.

[10]This Markov formulation is without loss of generality. It is well-known that any dynamic model can be described using Markov notation by defining large enough private states (see, e.g., Athey and Segal (2013) for a discussion).

[11]This, however, does not prevent one agent's reports from affecting future types of other agents, as well as the future public state through the public decisions.

decisions in a general way. In Section 6, we discuss extensions of the model with private actions.

Timing. At period $t = 0$, all the agents observe the initial state $\theta_0 = \hat{\theta}_0 \in \mathbf{\Theta}$, and a decision $x_0 \in \boldsymbol{X}$ is made. Then, in each period $t \geqslant 1$, every agent $i \in \mathcal{I}$ makes a (public) announcement $\hat{\theta}^i_t \in \mathbf{\Theta}^i$ of his private type, and all the agents observe the verifiable public state $\theta^0_t \in \mathbf{\Theta}^0$. Given the announced types and the public state, transfers y^i_t to each agent are calculated, and a decision $x_t \in \boldsymbol{X}$ is made.

Strategy. Let $\hat{\theta}_t = \left(\hat{\theta}^0_t, \hat{\theta}^1_t, ..., \hat{\theta}^N_t\right)$ denote the vector of reported types in period t including the public signal $\hat{\theta}^0_t = \theta^0_t$. The public history $h_t = (\hat{\theta}_0, \hat{\theta}_1, \ldots, \hat{\theta}_t, x_0, x_1, \ldots, x_t)$ contains all reports and public decisions[12] up to the period t, and we denote by $\mathcal{H}_t$ the set of all these public histories. A pure strategy of agent i is given by $s^i = \{s^i_t\}_{t=1}^T$, where

$$s^i_t : \left(\mathbf{\Theta}^{\boldsymbol{i}}\right)^t \times \mathcal{H}_{t-1} \to \mathbf{\Theta}^{\boldsymbol{i}} \tag{3}$$

maps agent i's private types up to period t and public histories before period t to a report in period t. We denote the set of all mixed strategies for player i by $\mathcal{S}^i$. We say that an agent follows a truth-telling strategy if $s^i_t \equiv \theta^i_t$ for every period $t \geqslant 1$.

Mechanism. A mechanism consists of the following.

1. For our results, it will be sufficient to define a **decision policy**[13] as a function

$$\mathsf{x} : \{0, 1, ..., T\} \times \mathbf{\Theta} \to \boldsymbol{X}, \tag{4}$$

   that determines the public decision $x_t = \mathsf{x}_t(\hat{\theta}_t)$ at every time period t.

2. A **transfer rule** $y^i_t = y^i_t(h_t)$ for every agent $i \in \mathcal{I}$, and every time period t, as a function of the public history.

   We say that a mechanism is *budget-balanced* if the sum of all transfers to all agents in every period t is zero:

$$\sum_{i \in \mathcal{I}} y^i_t \equiv 0 \ \forall t. \tag{5}$$

[12] The mechanism will use a deterministic decision policy, a deterministic function of the reports. Therefore, for simplicity, the public decisions will be omitted from the history.

[13] More general decision policies considering past reports (and randomization) are possible, see Bergemann and Välimäki (2019, Section 2.1). Due to the Markovian nature of the model, these cannot yield any higher social surplus.

4 Guaranteed Utility Mechanism (GUM)

To formally define a decision policy and a transfer rule in our mechanism, we introduce the following notation. We fix an ordering of the agents $\{0, 1, ..., N\}$ including "0" indicating the public signal. The intuition behind the following definition is that we consider the progress of the game as if the types were reported sequentially in this order in every period. We define the **history** of public reports up to $i \in \{0, 1, ..., N\}$ at period t as follows.

$$h_{t,i} = \big(\, \hat{\theta}_0, \hat{\theta}_1, ..., \hat{\theta}_{t-1}, (\hat{\theta}_t^0, \hat{\theta}_t^1, ..., \hat{\theta}_t^i)\, \big) \tag{6}$$

For a history $h = h_{t,i}$ let $\tilde{\theta} \sim \mu(h, \mathsf{x})$ denote the probabilistic continuation of the game $\tilde{\theta} = (\tilde{\theta}_0, \tilde{\theta}_1, ..., \tilde{\theta}_T)$.[14] When we take expectation on a random variable $\tilde{\theta}$ from $\mu(h, \mathsf{x})$, we will use the simplified notation $\mathbb{E}^{\mu(h,\mathsf{x})}\big[f(\tilde{\theta})\big] = \mathbb{E}^{\tilde{\theta}\sim\mu(h,\mathsf{x})}\big[f(\tilde{\theta})\big]$.

Decision policy. GUM chooses an *efficient* decision policy χ that maximizes the expected sum of all the agents' payoffs for the initial state $h_{0,N} = \hat{\theta}_0 = \theta_0 \in \boldsymbol{\Theta}$:

$$\chi \in \arg\max_{\chi'} \quad \mathbb{E}^{\mu(\theta_0,\chi')}\Bigg[\, \sum_{i\in\mathcal{I}} \sum_{t=0}^{T} u^i\big(\tilde{\theta}_t^i, \chi'(\tilde{\theta}_t)\big)\,\Bigg]. \tag{7}$$

Transfers. We define the **anticipated payoff** of agent j as

$$\Upsilon_{t,i}^j = \mathbb{E}^{\mu(h_{t,i},\chi)}\Big[\, \sum_{t'=0}^{T} u^j\big(\tilde{\theta}_{t'}^j, \chi(\tilde{\theta}_{t'})\big)\Big]. \tag{8}$$

Now we define the **externality** agent i's report $\hat{\theta}_t^i$ imposes on agent j as its total change in the anticipated payoff of j.

$$\gamma_t^{i\to j} = \Upsilon_{t,i}^j - \Upsilon_{t,i-1}^j = \sum_{t'=0}^{T}\Bigg(\mathbb{E}^{\mu(h_{t,i},\chi)}\Big[\,u^j\big(\tilde{\theta}_{t'}^j, \chi(\tilde{\theta}_{t'})\big)\Big] - \mathbb{E}^{\mu(h_{t,i-1},\chi)}\Big[\,u^j\big(\tilde{\theta}_{t'}^j, \chi(\tilde{\theta}_{t'})\big)\Big]\Bigg), \tag{9}$$

[14] Formally, we define it as follows. If $t' < t$ or ($t' = t$ and $j \le i$), then $\tilde{\theta}_{t'}^j = \hat{\theta}_{t'}^j$. Otherwise $\tilde{\theta}_{t'}^j$ is recursively defined by the stochastic type transition rule $\mu(\tilde{\theta}_{t'-1}^j, \mathsf{x}(\tilde{\theta}_{t'-1}))$.

and agent i pays a compensation $-\gamma_t^{i\to j}$ for his externality to j. To sum up, the **transfer** to agent i at period t is defined as

$$y_t^i = \sum_{j\neq i} \left(\gamma_t^{i\to j} - \gamma_t^{j\to i}\right) = \underbrace{\sum_{j\neq i} \gamma_t^{i\to j}}_{i\text{'s externalities on other agents}} - \underbrace{\sum_{j\neq i} \gamma_t^{j\to i}}_{\text{other agents' externalities on } i} . \tag{10}$$

The externality $\gamma_t^{i\to j}$ reflects the change in the anticipated payoff of agent j as a result of updating agent i's report at period t. The payments in GUM are defined so that for each report $\hat{\theta}_t^i$, each other agent j has to pay $\gamma_t^{i\to j}$ to agent i. When calculating the externalities caused by a report $\hat{\theta}_t^i$, we use the updated period-t reports for agents $1, ..., i-1$, whereas we keep period-$(t-1)$ reports for agents $i+1, ..., N$. We then calculate how the anticipated payoff of each other agent j changed. In other words, when calculating the effect of an agent's report, we update the newly reported types sequentially according to the order of the agents $0, 1, ..., N$, where 0 indicates the public signal.

In our mechanism, only an agent j who is *affected* positively (negatively) by a new report of agent i pays (gets) a corresponding transfer. For example, if agent i's report changes the expected total payoff of agent j by 15 and of agent k by -5, then j transfers 15 to i and k is paid 5 by i. Consequently, agent i collects $15 - 5 = 10$. The transfers in (10) are budget-balanced because every payment is from one agent to another.

We define the ***Guaranteed Utility Mechanism*** **(GUM)** to be a mechanism with an efficient decision policy, together with the balanced transfers defined in (10) that are constructed from the payments in (9). The name reflects the fact that it satisfies the following property.

4.1 The Guaranteed Utility Property (GUP)

We say that a mechanism satisfies the *Guaranteed Utility Property* (GUP) if there exists a strategy profile $s_* \in \mathcal{S}^{\mathcal{I}}$ and a vector $C \in \mathbb{R}^N$ such that:

$$\forall i \in \mathcal{I},\ \forall s^{-i} \in \mathcal{S}^{-i}: \qquad \mathbb{E}\left[U^i(s_*^i, s^{-i})\right] \geqslant C^i; \tag{11}$$

$$\sup_{s\in\mathcal{S}^{\mathcal{I}}} \sum_{i\in\mathcal{I}} \mathbb{E}\left[U^i(s)\right] = \sum_{i\in\mathcal{I}} C^i. \tag{12}$$

Intuitively, GUP requires that every agent i obtains an expected utility guarantee C^i by following the strategy s_*^i *regardless* of the other agents' strategies. In addition, GUP requires that the sum of all the utility guarantees equals the highest total expected surplus.

Equivalently, a mechanism satisfies the GUP property if and only if

$$\sum_{i\in\mathcal{I}} \max_{s^i\in\mathcal{S}^i} \min_{s^{-i}\in\mathcal{S}^{-i}} \mathbb{E}\big[U^i(s^i,s^{-i})\big] = \max_{s\in\mathcal{S}^{\mathcal{I}}} \sum_{i\in\mathcal{I}} \mathbb{E}\big[U^i(s)\big]. \tag{13}$$

In other words, the mechanism satisfies GUP if it induces a game in which the sum of all the agents' maximin expected utilities equals the highest total expected surplus.

GUM satisfies GUP. In GUM, each agent i guarantees an expected utility equal to his initial anticipated payoff $C^i = \Upsilon^i_{0,N} = \mathbb{E}^{\mu(\theta_0,\chi)}\big[\sum_{k=0}^{T} u^i\big(\tilde{\theta}^i_k, \chi(\tilde{\theta}_k)\big)\big]$ by being truthful. This outcome holds even if some other agents deviate from truth-telling.

Consider an arbitrary period t, and suppose that reports of all the agents $k < i$ in the order $\{1, ..., N\}$ have been updated. Then, the updated report $\hat{\theta}^i_t$ of an agent i changes the anticipated payoff of agent j by $\Upsilon^j_{t,i} - \Upsilon^j_{t,i-1}$, and i compensates j the amount of $-\gamma^{i\to j}_t = \Upsilon^j_{t,i-1} - \Upsilon^j_{t,i}$. For the martingale property, it is sufficient to prove that, in this bilateral interaction between i and j, if either agent is truthful, then he is unaffected by the dishonesty of the other agent in expectation. Consider agent j. If j is truthful, then the transfer from i to j compensates for the change in the anticipated expected total payoff of j, regardless of the true types of i and the other agents, $-\gamma^{i\to j}_t = \Upsilon^j_{t,i-1} - \Upsilon^j_{t,i}$. Consider agent i. If i is truthful, then the report $\hat{\theta}^i_t$ is distributed according to the $\mu\big(\hat{\theta}^i_{t-1}, \chi(\hat{\theta}_{t-1})\big)$. Thus, regardless of j being truthful or not, the law of iterated expectations implies that the expected change in the anticipated payoff on j must be zero, $\mathbb{E}[\gamma^{i\to j}_t] = \mathbb{E}[\Upsilon^j_{t,i}] - \mathbb{E}[\Upsilon^j_{t,i-1}] = 0.$[15] Therefore, the expected compensation paid by a truthful agent i is zero.

Given the martingale property, the final anticipated payoff of every honest agent i equals his payoff $\Upsilon^i_{T,N} = \sum_{k=0}^{T} u^i\big(\tilde{\theta}^i_k, \chi(\tilde{\theta}_k)\big)$. As a result, the final value of this martingale equals i's total utility, $\mathbb{E}[\Upsilon^i_{T,N}] = \mathbb{E}[U^i]$, which implies (11) with $C^i = \Upsilon^i_{0,N}$. If every agent is truthful, then the mechanism implements an efficient decision policy. As a result, the sum of all the utility guarantees equals the highest total expected surplus, $\sum_{i\in\mathcal{I}} \mathbb{E}\big[U^i(s_*)\big] = \sum_{i\in\mathcal{I}} \mathbb{E}^{\mu(\theta_0,\chi)}\Big[\sum_{k=0}^{T} u^i\big(\tilde{\theta}^i_k, \chi(\tilde{\theta}_k)\big)\Big] = \sum_{i\in\mathcal{I}} \Upsilon^i_{0,N} = \sum_{i\in\mathcal{I}} C^i$, which implies (12).

A direct consequence of GUM satisfying GUP is that GUM sustains truth-telling strategies as a Bayesian Nash equilibrium. Moreover, in every BNE sustained by GUM, every agent

[15]Formally, $\mathbb{E}[\Upsilon^j_{t,i} \mid h_{t,i-1};\ \hat{\theta}^i_t = \theta^i_t] = \mathbb{E}^{\hat{\theta}^i_t\sim\mu(\hat{\theta}^i_{t-1},\chi(\hat{\theta}_{t-1}))}\Big[\sum_{t'=0}^{T} \mathbb{E}^{\tilde{\theta}\sim\mu((h_{t,i-1},\hat{\theta}^i_t),\chi)}\Big[u^j\big(\tilde{\theta}^j_{t'}, \chi(\tilde{\theta}_{t'})\big)\Big]\Big]$

$$= \sum_{t'=0}^{T} \mathbb{E}^{\mu(h_{t,i-1},\chi)}\Big[u^j\big(\tilde{\theta}^j_{t'}, \chi(\tilde{\theta}_{t'})\big)\Big] = \mathbb{E}\big[\Upsilon^j_{t,i-1}\big].$$

obtains the same expected utility equal to his initial anticipated payoff assuming truthful reports and an efficient decision policy. For instance, suppose a strategy profile $s \in \mathcal{S}^{\mathcal{I}}$ is a Bayesian Nash equilibrium and s^i_* is the truthful strategy. On the one hand, the expected utility of agent i is bounded from below by the utility guarantee of $\Upsilon^i_{0,N}$:

$$\mathbb{E}\big[U^i(s)\big] \overset{BNE}{\geqslant} \mathbb{E}\big[U^i(s^i_*, s^{-i})\big] \overset{GUP}{\geqslant} \Upsilon^i_{0,N}, \tag{14}$$

where the first inequality follows from the definition of BNE, and the second inequality follows from the lower bound on the agent's utility being the utility guarantee of $\Upsilon^i_{0,N}$. On the other hand, the expected utility of agent i is bounded from above by $\Upsilon^i_{0,N}$:

$$\begin{aligned}\mathbb{E}\big[U^i(s)\big] &= \mathbb{E}\Big[\sum_j U^j(s)\Big] - \sum_{j\neq i}\mathbb{E}\big[U^j(s)\big] \\ &\leqslant \sum_{j\in\mathcal{I}}\Upsilon^j_{0,N} - \sum_{j\neq i}\mathbb{E}\big[U^j(s)\big] \overset{(14)}{\leqslant} \sum_j \Upsilon^j_{0,N} - \sum_{j\neq i}\Upsilon^j_{0,N} = \Upsilon^i_{0,N},\end{aligned} \tag{15}$$

where the first inequality holds because $\sum_{j\in\mathcal{I}}\Upsilon^j_{0,N}$ is the highest possible total expected utility, and the second follows from (14) since $\sum_{j\neq i}\mathbb{E}\big[U^j(s)\big] \geqslant \sum_{j\neq i}\Upsilon^j_{0,N}$. Therefore, in every Bayesian Nash equilibrium, agent i's expected utility is given by $\Upsilon^i_{0,N}$. We summarize the results below:[16]

Proposition 1. *All BNE sustained by GUM are efficient and utility-equivalent.*

4.2 Collusion-Proofness

We introduce a collusive environment that allows the agents to jointly coordinate their reported types and to make balanced transfers within the colluding coalition. We assume that the reports and the transfers are enforced by a (benevolent) coordinator who has all the information the agents in the coalition have.[17] Specifically, for a set of agents $L \subseteq \mathcal{I}$, a side contract is determined by $\bar{s}^L = \{\bar{s}^L_t\}_{t=1}^T$, where

$$\bar{s}^L_t : \big(\boldsymbol{\Theta}^{\boldsymbol{L}}\big)^t \times \mathcal{H}_t \to \boldsymbol{\Theta}^{\boldsymbol{L}} \times \mathbb{R}^{|L|}. \tag{16}$$

A pure collusive strategy $\bar{s}^L$ maps the joint information of the agents in L to reports $\hat{\theta}^L$ and budget-balanced transfers within the coalition. We denote by $\bar{\mathcal{S}}^L$ the set of probability

[16] The formal proof of Proposition 1 is omitted as it directly follows from Theorem 1 proven in Section 4.2.

[17] It is without loss of generality to assume that collusion is enforced by a side contract (see Laffont and Martimort (1997) and Laffont and Martimort (2000)).

distributions on these pure strategies. Intuitively, a side contract determines coordination on (possibly random) submitted reports and budget-balanced transfers within L.

We denote by $\bar{s}^{\mathcal{I}\setminus\{i\}} \in \bar{S}^{\mathcal{I}\setminus\{i\}}$ a joint collusive strategy of all the agents except i. The collusive strategy $\bar{s}^{\mathcal{I}\setminus\{i\}}$ can be treated as a "punishment" the other agents can coordinate against agent i who does not join the grand coalition.

We say that a collusive strategy $\bar{s}^{\mathcal{I}} \in \bar{\mathcal{S}}^{\mathcal{I}}$ of the grand coalition is a ***weak collusive equilibrium*** if:[18]

$$\forall i \in \mathcal{I},\ \forall s^i \in \mathcal{S}^i,\ \exists \bar{s}^{\mathcal{I}\setminus\{i\}} \in \bar{\mathcal{S}}^{\mathcal{I}\setminus\{i\}}: \qquad \mathbb{E}\big[U^i(s^i, \bar{s}^{\mathcal{I}\setminus\{i\}})\big] \leqslant \mathbb{E}\big[U^i(\bar{s}^{\mathcal{I}})\big]. \tag{17}$$

Intuitively, condition (17) states that the grand coalition collusive strategy $\bar{s}^{\mathcal{I}}$ can be sustained if no agent i is able to deviate and collect a strictly higher utility against the "punishment" strategy $\bar{s}^{\mathcal{I}\setminus\{i\}}$ all other agents can impose on him. The right-hand side of (17) reflects i's utility if he follows the grand coalition collusive strategy. The left-hand side of (17) reflects i's utility if he is not part of the grand coalition, i.e. if he follows a non-collusive strategy. Thus, $\bar{s}^{\mathcal{I}}$ is a weak collusive equilibrium if, for each agent i, there exists a "threat" $\bar{s}^{\mathcal{I}\setminus\{i\}}$ all other agents can use to incentivize agent i to join the grand coalition. A weak collusive equilibrium should be interpreted as a necessary but not a sufficient condition for a collusive strategy profile to be plausible. For instance, every Bayesian Nash equilibrium satisfies the definition of weak collusive equilibrium.

An alternative way to define a weak collusive equilibrium $\bar{s}^{\mathcal{I}}$ is to require that no agent has a strictly dominant deviation, where a *strictly dominant deviation from* $\bar{s}^{\mathcal{I}}$ is a strategy s^i such that:

$$\forall \bar{s}^{\mathcal{I}\setminus\{i\}} \in \bar{\mathcal{S}}^{\mathcal{I}\setminus\{i\}}: \qquad \mathbb{E}\big[U^i(s^i, \bar{s}^{\mathcal{I}\setminus\{i\}})\big] > \mathbb{E}\big[U^i(\bar{s})\big]. \tag{18}$$

We now introduce our notion of collusion-proofness.

We say that a game is ***collusion-proof*** if all weak collusive equilibria are utility-equivalent. That is, the vector of agents' utilities $\Big(\mathbb{E}\big[U^i(\bar{s}^{\mathcal{I}})\big]\Big)_{i\in\mathcal{I}}$ is the same for all weak collusive equilibria $\bar{s}^{\mathcal{I}}$.

We will prove that this definition of collusion-proofness implies not only the guaranteed utility property but also a stronger version of the guaranteed utility property.

Theorem 1. *The Guaranteed Utility Mechanism is collusion-proof.*
Proof: See Appendix B.

[18] The order of $\forall s^i \in \mathcal{S}^i$ and $\exists \bar{s}^{\mathcal{I}\setminus\{i\}} \in \bar{\mathcal{S}}^{\mathcal{I}\setminus\{i\}}$ in (17) are interchangeable since it is an instance of a 2-player (agent i and coalition $\mathcal{I} \setminus \{i\}$) 0-sum game.

The reason GUM is collusion-proof is that no coalition of agents $L \subseteq \mathcal{I}$ can collude and collect a total utility higher than $\sum_{j\in L} \Upsilon^j_{0,N}$ if other non-colluding agents are truthful. For any coalition of agents $L \subseteq \mathcal{I}$, the upper bound for the sum of their expected utilities is the highest total expected surplus of all agents minus the sum of the guaranteed expected utilities of other agents $i \notin L$ outside of the coalition. If members $\mathcal{I} \setminus L$ outside of the coalition follow the equilibrium truth-telling strategies, coalition L cannot get a higher total expected utility:

$$\begin{aligned}\sup_{\tilde{\theta}} \sum_{i\in\mathcal{I}} \mathbb{E}^{\mu(\theta_0,\chi)}\Big[U^i(\theta^i, \chi(\tilde{\theta}), y^i)\Big] - \sum_{i\notin L} \mathbb{E}^{\mu(\theta_0,\chi)}\Big[U^i(\theta^i, \chi(\hat{\theta}), y^i)\Big] \\ \leqslant \sum_{i\in\mathcal{I}} \Upsilon^i_{0,N} - \sum_{i\notin L} \Upsilon^i_{0,N} = \sum_{j\in L} \Upsilon^j_{0,N}.\end{aligned} \tag{19}$$

To highlight the connection between utility guarantees and collusion-proofness, we introduce the following definition that extends the GUP property to a collusive environment. We say that a mechanism satisfies the ***Collusion-Proof Guaranteed Utility Property*** (CPGUP) if there exists a strategy profile $s_* \in \mathcal{S}^{\mathcal{I}}$ and a vector $C \in \mathbb{R}^N$ satisfying the following:

$$\forall i \in \mathcal{I},\ \forall \bar{s}^{\mathcal{I}\setminus\{i\}} \in \bar{\mathcal{S}}^{\mathcal{I}\setminus\{i\}}: \qquad \mathbb{E}\big[U^i(s^i_*, \bar{s}^{\mathcal{I}\setminus\{i\}})\big] \geqslant C^i; \tag{20}$$

$$\sup_{\bar{s}^{\mathcal{I}}\in\bar{\mathcal{S}}^{\mathcal{I}}} \sum_{i\in\mathcal{I}} \mathbb{E}\big[U^i(\bar{s}^{\mathcal{I}})\big] = \sum_{i\in\mathcal{I}} C^i. \tag{21}$$

Equivalently, a mechanism satisfies the CPGUP property if and only if

$$\sum_{i\in\mathcal{I}} \max_{s^i\in\mathcal{S}^i} \min_{\bar{s}^{\mathcal{I}\setminus\{i\}}\in\bar{\mathcal{S}}^{\mathcal{I}\setminus\{i\}}} \mathbb{E}\big[U^i(s^i, \bar{s}^{\mathcal{I}\setminus\{i\}})\big] = \max_{\bar{s}^{\mathcal{I}}\in\bar{\mathcal{S}}^{\mathcal{I}}} \sum_{i\in\mathcal{I}} \mathbb{E}\big[U^i(\bar{s}^{\mathcal{I}})\big]. \tag{22}$$

We now discuss the connection between the CPGUP property and GUP introduced in Section 4.1. Recall that, in (11) and (12), agent i evaluates his utility against the non-collusive strategies of other agents. That is, GUP property requires the existence of a utility guarantee for agent i, assuming the other agents do not coordinate their strategies. However, the strong version of GUP requires that a utility guarantee exists for agent i even if all the other agents collude and coordinate their reports. As can be seen, agent i's utility on the left-hand side of (20) is evaluated against any collusive strategy $\bar{s}^{\mathcal{I}\setminus\{i\}}$ of agents $\mathcal{I} \setminus \{i\}$. Similarly, the left-hand side of (21) allows any collusive strategy $\bar{s}^{\mathcal{I}}$ of the grand coalition.

Proposition 2. *A mechanism is collusion-proof if and only if it satisfies CPGUP.*

Proof. On the one hand, if G satisfies CPGUP, then

$$C^i \overset{(11)}{\leq} \mathbb{E}\big[U^i(s_*^i, \bar{s}^{\mathcal{I}\setminus\{i\}})\big] \overset{(17)}{\leq} \mathbb{E}\big[U^i(\bar{s}^{\mathcal{I}})\big], \tag{23}$$

Thus,

$$\sum_{i\in\mathcal{I}} C^i \overset{(23)}{\leq} \sum_{i\in\mathcal{I}} \mathbb{E}\big[U^i(\bar{s}^{\mathcal{I}})\big] \overset{(12)}{\leq} \sum_{i\in\mathcal{I}} C^i, \tag{24}$$

therefore, $\mathbb{E}\big[U^i(\bar{s}^{\mathcal{I}})\big] = C^i$ for every agent $i \in \mathcal{I}$.

On the other hand, if G does not satisfy CPGUP, then the negation of (22) says that there exists $(\bar{s}^{\mathcal{I}}, \bar{s}^{\mathcal{I}\setminus\{1\}}, \bar{s}^{\mathcal{I}\setminus\{2\}}, ..., \bar{s}^{\mathcal{I}\setminus\{N\}})$ satisfying the following.

$$\sum_{i\in\mathcal{I}} \max_{s^i\in\mathcal{S}^i} \mathbb{E}\big[U^i(s^i, \bar{s}^{\mathcal{I}\setminus\{i\}})\big] < \sum_{i\in\mathcal{I}} \mathbb{E}\big[U^i(\bar{s}^{\mathcal{I}})\big]. \tag{25}$$

We can modify the transfer rules in $\bar{s}^{\mathcal{I}}$ by an additional constant balanced vector $c^{\mathcal{I}}$, denoted by $\bar{s}^{\mathcal{I}} + c^{\mathcal{I}}$. Any profile of expected utilities $(U_*^1, U_*^2, ..., U_*^N)$ satisfying $\sum_{i\in\mathcal{I}} \mathbb{E}\big[U^i(\bar{s}^{\mathcal{I}})\big] = \sum_{i\in\mathcal{I}} U_*^i$ can be achieved by a collusive strategy profile of the form $\bar{s}^{\mathcal{I}} + c^{\mathcal{I}}$. There are different expected utility vectors which also satisfy $\max_{s^i\in\mathcal{S}^i} \mathbb{E}\big[U^i(s^i, \bar{s}^{\mathcal{I}\setminus\{i\}})\big] \leq U_*^i$ for every agent $i \in \mathcal{I}$. For all these vectors, $\bar{s}^{\mathcal{I}} + c^{\mathcal{I}}$ is a weak collusive equilibrium, witnessed by $\big(\bar{s}^{\mathcal{I}\setminus\{1\}}, \bar{s}^{\mathcal{I}\setminus\{2\}}, ..., \bar{s}^{\mathcal{I}\setminus\{N\}}\big)$. □

5 Relation to Bergemann and Välimäki (2010) and Athey and Segal (2013)

We now discuss GUM in relation to the Dynamic Pivot Mechanism (DPM) and the Balanced Team Mechanism (BTM). To begin, we extend the definition of $\gamma^{i\to j}$ in (9) as follows:

$$\gamma_t^{i\to j}(h) = \sum_{t'=0}^{T} \left(\mathbb{E}^{\mu((h,\hat{\theta}_t^i),\chi)}\Big[u^j\big(\tilde{\theta}_{t'}^j, \chi(\tilde{\theta}_{t'})\big)\Big] - \mathbb{E}^{\mu(h,\chi)}\Big[u^j\big(\tilde{\theta}_{t'}^j, \chi(\tilde{\theta}_{t'})\big)\Big] \right). \tag{26}$$

This extended definition will allow us to illustrate better the technical differences among the three mechanisms.[19] In particular, the payment rules in the three mechanisms can be expressed as follows:

[19] The main properties of the three mechanisms (DPM, BTM, and GUM) were summarized in Table 1.

$$y_t^{DPM,i} = \sum_{j\neq i} \gamma_t^{i\to j}(h_{t,\mathcal{I}\setminus\{i\}}), \tag{27}$$

$$y_t^{BTM,i} = \sum_{j\neq i} \gamma_t^{i\to j}(h_{t-1,N}) - \frac{1}{N-1}\sum_{j\neq i}\sum_{k\neq j} \gamma_t^{j\to k}(h_{t-1,N}), \tag{28}$$

$$y_t^{GUM,i} = \sum_{j\neq i} \Big(\gamma_t^{i\to j}(h_{t,i-1}) - \gamma_t^{j\to i}(h_{t,j-1})\Big), \tag{29}$$

where $h_{t,\mathcal{I}\setminus\{i\}} = (\hat{\theta}_0, \hat{\theta}_1, ..., \hat{\theta}_{t-1}, \hat{\theta}_t^{-i})$ is the history of reports until time point t excluding $\hat{\theta}_t^i$.

Additive externality. As can be seen from (27), (28), and (29), the key differences among the three mechanisms (DPM, BTM, and GUM) are the way the externalities and the contributions of an individual report to the externalities are calculated. In our mechanism, these contributions are additive. That is, when multiple reports are made, the joint externality imposed by simultaneously updating these reports equals the sum of the calculated individual externalities. Moreover, each agent i pays the transfer to agent j equal to the externality j imposes on i. If the externality is not additive, then members of a collusive cartel can exploit the difference between the sum of their calculated individual externalities and the true joint effect of simultaneous updating of their reports. This makes a mechanism prone to collusion.

We now illustrate formally that the way the externality is defined in GUM additive. That is, the sum of all the agents' externalities is equal to the total externality imposed by updating all the agents' reports simultaneously.

We define the ***total* externality** *simultaneous* updating of all the agent's reports at period t imposes on the agent's j payoff as follows:

$$\gamma_t^{\mathcal{I}\to j} = \sum_{t'=t}^{T} \left(\mathbb{E}^{\mu(h_{t,N},\chi)}\left[u^j\left(\tilde{\theta}_{t'}^j, \chi(\tilde{\theta}_{t'})\right)\right] - \mathbb{E}^{\mu(h_{t-1,N},\chi)}\left[u^j\left(\tilde{\theta}_{t'}^j, \chi(\tilde{\theta}_{t'})\right)\right] \right). \tag{30}$$

That is, $\gamma_t^{\mathcal{I}\to j}$ reflects the change in the j's anticipated payoff as a result of updating all the agents' reports at period t. In other words, we calculate how the anticipated payoff of agent j changes when we update the reported types from $\hat{\theta}_{t-1,N}$ to $\hat{\theta}_{t,N}$. Intuitively, we update the newly reported types simultaneously when calculating the total externality.

The externality is ***additive*** if the sum of the externalities across all the agents is equal to the total externality:

$$\sum_{i\in\mathcal{I}} \gamma_t^{i\to j} = \gamma_t^{\mathcal{I}\to j} \quad \forall j,\ \forall t. \tag{31}$$

Lemma 1. *The externality defined in* (9) *is additive.*
Proof: See Appendix A.

Note that (9) uses the total externality defined as the simultaneous updating of all the agent's reports. More generally, we can define the externality to be additive if the simultaneous updating of reports of a coalition of agents $L \subseteq \mathcal{I}$ equals the sum of the externalities across the agents in that coalition. GUM satisfies this definition of additivity as well.

If the externality is additive, the change in every agent's anticipated payoff resulting from updating types from $\hat{\theta}_{t-1,N}$ to $\hat{\theta}_{t,N}$ at period t could be decomposed into the sum of the marginal "contributions" of every agent. This additivity is crucial for the mechanism to have the Collusion-Proof Guaranteed Utility Property. Contrarily, if the externality is not additive, then it provides an opportunity for several agents to jointly deviate from truthtelling and collect a higher utility. For instance, this is the case with the Balanced Team Mechanism.

5.1 Examples for AGV (static BTM) and BTM

We now provide an example in which the truth-telling strategy sustained by the BTM is weakly dominated. Moreover, only a non-truthful inefficient equilibrium survives the iterated elimination of weakly dominated strategies (regardless of the order of elimination).

Static case: AGV (which coincides with BTM for $T = 1$). This weakness (truth-telling strategy being weakly dominated and the fact that only a non-truthful equilibrium survives the iterated elimination of weakly dominated strategies) of BTM is already present in the single-round setup, where BTM coincides with the AGV mechanism.

If there is a type profile for which there are multiple efficient decisions, then there are multiple efficient decision policies. We will call an AGV with a particular efficient decision policy an AGV implementation. We state the results below, and the proofs are in Appendix E.

Theorem 2. *There exists a setup with multiple efficient decision policies, such that, in one of its AGV implementations, the truthful strategy of an agent is weakly dominated. Moreover, an inefficient strategy profile survives iterative elimination of weakly dominated strategies (not depending on the order of elimination).*
Proof: See Appendix E.

To highlight the role of multiple efficient decision policies in our results, we provide the following Theorem.

Theorem 3. *If a setup has a unique efficient decision policy, then under the unique AGV implementation, the truthful strategy profile survives iterative elimination of weakly dominated strategies.*
Proof: See Appendix E.

Dynamic case: BTM ($T > 1$). We now illustrate that Theorem 3 no longer holds in a dynamic environment.

Theorem 4. *There exists a setup with a unique efficient decision policy such that, in the unique BTM implementation, only an inefficient strategy profile survives the iterative elimination of weakly dominated strategies (not depending on the order of elimination).*
Proof: See Appendix D.

Although the formal proof of Theorem 4 is delegated to Appendix D, we now illustrate the main idea using Example 5.1 below.

Example 5.1. *Consider a game with a large number of periods T and with three agents: blue (b), red (r), and green (g). The set of agents is denoted by $\mathcal{I} = \{b, r, g\}$. At periods 1, 2, ..., $T-1$, nothing happens except that the types of blue and red stochastically evolve, independently. At period T, the type of both blue and red is either H (high type) or L (low type). A public decision "YES" or "NO" is made. If the decision is YES, then the payoffs are $u_T^i(L,\ YES) = 5$ and $u_T^i(H,\ YES) = 100$ for $i \in \{b, r\}$, and $u_2^g(YES) = -150$. If the decision is NO then all the agents' payoffs are normalized to zero. Thus, it is efficient to choose YES if both blue and red agents have type H, and NO otherwise.*

For $i \in \{b, r\}$, let $p_t^i = \Pr(\theta_T^i = H \mid \theta_t^i)$ denote the probability at t that i will have type H in round T. The transfers can be directly expressed by these probabilities.

The Balanced Team Mechanism. Each agent publicly reports the probability $\hat{p}_t^i \in [0, 1]$ at every period $t = 1, 2, ..., T$, and

$$\hat{p}_0^i = p_0^i = \mathbb{E}(p_t^i). \tag{32}$$

Given the reports, each agent gets a transfer calculated in two steps. First, for each t, we calculate how each agent's report changes the anticipated payoffs of other agents, given all the reported types in the previous round $t-1$:

$$\gamma_t^{BTM,i} = \sum_{j\in\mathcal{I}\setminus\{i\}} \sum_{t'=0}^{T} \mathbb{E}\Big[u_{t'}^j\big(\hat{p}_t^i,\ \hat{p}_{t-1}^{\mathcal{I}\setminus\{i\}}\big) - u_{t'}^j\big(\hat{p}_{t-1}^{\mathcal{I}}\big)\Big], \tag{33}$$

where $u_{t'}^j$ expresses the payoff of agent j in round t', and $\hat{p}_{t-1}^{\mathcal{I}}$ denotes the reports of all agents $\mathcal{I} = \{b, r, g\}$ at $t-1$.

Second, each agent i gets this transfer $\gamma_t^{BTM,i}$ paid equally by the other two agents. Formally, the total transfer to agent i becomes:

$$y^{BTM,i} = \sum_{t=1}^{T} \Big[\gamma_t^{BTM,i} - \frac{1}{2} \sum_{j\in\mathcal{I}\setminus\{i\}} \gamma_t^{BTM,j}\Big]. \tag{34}$$

In this example, each agent gets a transfer equal to the imposed externality on others,

$$\gamma_t^{BTM,i} = 50 \cdot \hat{p}_{t-1}^{-i}(\hat{p}_{t-1}^{i} - \hat{p}_t^{i}), \tag{35}$$

because the report changed the probability of YES from $\hat{p}_{t-1}^{-i} \cdot \hat{p}_{t-1}^{i}$ to $\hat{p}_{t-1}^{-i} \cdot \hat{p}_t^{i}$, and this decision would cause a total of $100 - 150 = -50$ change in the total payoff of the other agents. Each of these transfers $\gamma_t^{BTM,i}$ is paid by the other two agents, shared equally.

Consider the following deviation from the truth-telling strategies. In every even round both blue and red agents report $\hat{p}_{2k}^b = \hat{p}_{2k}^r = 0$, and in every odd round they report $\hat{p}_{2k+1}^b = \hat{p}_{2k+1}^r = 1$. Then, from round $t = 2$ on, in every even round both get a transfer of $\gamma_{2k}^{BTM,i} = 50 \cdot 1 \cdot (1-0) = 50$, and in every odd round both get a transfer $\gamma_{2k+1}^{BTM,i} = 50 \cdot 0 \cdot (0-1) = 0$. So this deviation from the truthful strategies is beneficial for both agents.[20]

Note that, in this dynamic setup, deviation from truthfulness does not require coordination: if either agent starts playing one of the stable equilibria, then the other agent is incentivized to join. In Appendix D, we show a discrete example based on the same idea with a full analysis that iterative elimination of weakly dominated strategies eliminate all efficient strategy profiles.

Technical Differences: BTM vs. GUM. We now discuss the key technical differences between the BTM and the GUM in detail. In particular, we discuss the two technical differences between the way transfers are calculated in the BTM and the GUM.

[20] The described deviation from the truth-telling equilibrium is reminiscent of the stable (pure strategies) equilibrium in the Battle of Sexes. The truth-telling equilibrium is analogous to the unstable equilibrium.

Incentive Term. In the BTM, the externality the agent i's updated report $\hat{\theta}_t^i$ imposes on an agent j's payoff is given by:

$$\gamma_t^{BTM,i\to j} = \mathbb{E}^{\mu(\hat{\theta}_t^i, \hat{\theta}_{t-1}^{-i}, \chi)}\left[\sum_{t'} u^j\left(\theta_{t'}, \chi(\theta_{t'})\right)\right] - \mathbb{E}^{\mu(\hat{\theta}_{t-1}, \chi)}\left[\sum_{t'} u^j\left(\theta_{t'}^j, \chi(\theta_{t'})\right)\right]. \tag{36}$$

When calculating the change in the j's anticipated payoff at period t' as a result of updating agent i's report at period t, the BTM updates only the report of agent i from $\hat{\theta}_{t-1}^i$ to $\hat{\theta}_t^i$ while keeping reports of all other agents $j \neq i$ at period-$(t-1)$ value. That is, in both the anticipated payoff functions, we have $\hat{\theta}_{t-1}^{-i}$. This way of capturing the externality the agent i's report $\hat{\theta}_t^i$ imposes on other agents is not, in general, additive:[21]

$$\sum_{i\in\mathcal{I}} \gamma_t^{BTM,i\to j} \neq \gamma_t^{\mathcal{I}\to j}. \tag{37}$$

Intuitively, if the joint externality of several agents differs from the sum of their externalities, they can "game" the mechanism. That is, by misreporting their types, agents can exploit the difference between the total and the sum of the externalities their reports impose on other agents.

How transfers are paid. Second, the GUM differs from the BTM in the way transfers are paid. In the BTM, to ensure the transfers are budget-balanced, the agent i's incentive term is paid equally by all the other agents. In contrast, in the GUM, after each updated report θ_t^i, each other agent j pays the transfer equal to the effect of this report on j.

We illustrate the two differences related to the transfers using the motivating example with agents $\{b, r, g\}$ introduced above.

First, recall that, in the BTM, the payment $\gamma_t^{BTM,r}$ that reflects how the anticipated payoffs of blue and green change as a result of updating red's report from $\hat{p}_{t-1}^r$ to $\hat{p}_t^r$ is:

$$\gamma_t^{BTM,r} = 50\, \hat{p}_{t-1}^b \left(\hat{p}_{t-1}^r - \hat{p}_t^r\right),$$

and if $\hat{p}_{t-1}^b = 0$ then $\gamma_t^{BTM,r} = 0$ regardless of $\hat{p}_t^r$. This is the reason red can benefit from an exaggerated report $\hat{p}_t^r = 1$.

However, in the GUM, according to the order of agents $\{b, r, g\}$, we have

$$\gamma_t^{GUM,b} = 50\, \hat{p}_{t-1}^r \left(\hat{p}_{t-1}^b - \hat{p}_t^b\right) = \gamma_t^{BTM,b} \tag{38}$$

$$\gamma_t^{GUM,r} = 50\, \hat{p}_t^b \left(\hat{p}_{t-1}^r - \hat{p}_t^r\right) \not\equiv \gamma_t^{BTM,r}, \tag{39}$$

[21]If in every round, only one agent has randomness in his new type, then the externality is additive in the BTM.

because we already updated the report of red to $\hat{p}_t^r$ when calculating the incentive term for red. Note that with our way of calculating $\gamma_t^{GUM,r}$, agent red cannot report $\hat{p}_t^r = 1$ without "paying the price" as now he has to compensate for the externality his exaggerated report $\hat{p}_2^t = 1$ imposes on the other two agents.

Second, to guarantee that the mechanism is balanced, $\gamma_2^{GUM,r}$ is paid by the agents affected by $\hat{p}_2^r$. For instance, since green's payoff is -150 in case decision *YES* is made, red has to compensate green if his report $\hat{p}_2^r$ increases the chances of decision *YES* being made.

Note that only modifying the way the incentive terms are calculated is not sufficient. For example, suppose we use the modified value of γ_t^r but require it to be paid equally by all the other agents like in the BTM. Then, although red now must compensate for the externality his exaggerated report imposes on the other agents, he does not compensate for the entire change in the anticipated payoff of others. Thus his incentives to misreport remain. Simultaneously changing both the way the incentive terms are calculated and the way they are paid guarantees the truth-telling strategy is not weakly dominated.

5.2 Example (*Dynamic Pivot Mechanism*)

We now illustrate that the dynamic pivot mechanism introduced in Bergemann and Välimäki (2010) is subject to the possibility of inefficient equilibria. Furthermore, the truthful strategy sustained by the mechanism might be weakly dominated, and the iterated elimination of dominated strategies may eliminate all efficient outcomes.

Example. Consider a two-period game with three agents: blue (b), red (r), and green (g), and we denote the set of all agents by $\mathcal{I} = \{b, r, g\}$. In both periods, an allocation (i.e., a public decision) $a_t \in A$ is an element of a finite set $A = \{-1, 0, +1\}$ of possible allocations. The payoff function for all $i \in \mathcal{I}$ is defined as follows:

$$u_i(a_t, \theta_{i,t}) = \theta_{i,t} \cdot a_t - (a_t)^2, \tag{40}$$

where $\theta_{i,t}$ is the type of player i at period t, described next.

At $t = 1$, $\theta_{b,1} = \theta_{r,1} = 5$ and $\theta_{g,1} = -2$. At $t = 2$, the types depend on the types at $t = 1$ as well as the allocation a_1. If $a_1 = -1$, then $\theta_{b,2} = \theta_{r,2} = 0$; otherwise, $\theta_{b,2} = 100$ and $\theta_{r,2} = -100$. The green type at $t = 2$ is $\theta_{g,2} = 0$, regardless of the allocation. In addition, agents can report the outside option X, which represents the "efficient exit condition".[22] Intuitively, the agent is deemed irrelevant if he does not affect the efficient social decision

[22] As mentioned in Bergemann and Välimäki (2010), page 781, paragraph 2, X is a type with $u(X, a_t) = 0$ and the transition from (X, a_t) to type X occurs with probability one.

upon exit, the mechanism satisfies the efficient exit condition if agents neither make nor receive transfers in periods where they are irrelevant.

The Dynamic Pivot Mechanism. We now describe the dynamic pivot mechanism. Given the described types, the socially efficient allocation for every t is the following:

$$a_t = \text{sign}\left(\theta_{b,t} + \theta_{r,t} + \theta_{g,t}\right). \tag{41}$$

We now specify the non-zero monetary transfers as functions of the reported types $\hat{\theta}_{i,t}$. If $\hat{\theta}_{b,2} = 100$ and $\hat{\theta}_{r,2} = -100$, then both blue and red pay 99 at $t = 2$, $p_{b,2} = p_{r,2} = 99$. At $t = 1$, if $\hat{\theta}_{b,1} = X$, $\hat{\theta}_{r,1} = 5$, and $\hat{\theta}_{g,1} = -2$ then red pays 2 at $t = 1$, $p_{r,1} = 2$. Finally, if $\hat{\theta}_{r,1} = X$, $\hat{\theta}_{b,1} = 5$, and $\hat{\theta}_{g,1} = -2$ then blue pays 2 at $t = 1$, $p_{b,1} = 2$. The expected utilities are summarized in Table 3 below.

$\mathbb{E}U^b, \mathbb{E}U^r$	$\hat{\theta}_{r,1} = 5$	$\hat{\theta}_{r,1} = X$
$\hat{\theta}_{b,1} = 5$	$-94, -94$	$-96, -94$
$\hat{\theta}_{b,1} = X$	$-94, -96$	$-2, -2$

Table 2: utilities from the truthful and "exit" report for blue and red agents.

The main idea behind the example is that blue and red are better off reporting type X at $t = 1$ and then rejoining the mechanism by reporting their type truthfully at $t = 2$. By doing this, blue and red avoid the conflict as they want the opposite allocations at $t = 2$ (blue benefits from $a_2 = 1$ but red benefits from $a_2 = -1$). To see, note that, in both periods, green is better off being truthful as the truthful strategy dominates reporting $\hat{\theta}_{g,t} = X$). In addition, if the decision at $t = 1$ is $a_1 = 1$, then both blue and red are better off reporting $\hat{\theta}_{b,2} = 100$ and $\hat{\theta}_{r,2} = -100$ than reporting $\hat{\theta}_{b,2} = \hat{\theta}_{r,2} = X$.

6 Additional Properties and Extensions

We now discuss some of the additional properties and extensions of our mechanism.[23] First, in Section 6.1, we enhance GUM to satisfy a dynamic version of an interim participation constraint. Second, in Section 6.2 we extend our results to a more general environment where agents observe informative signals of the true past types of other agents. In Section 6.3, we allow agents to choose private actions.

[23]In Appendix H, we discuss a symmetric version of GUM and the connection to Shapley values.

6.1 Participation constraints

To focus on efficiency and budget balance, we presented the main model without explicitly addressing the agents' participation incentives. One interpretation of that setup is that the decisions and payments are implemented by an external enforcer. To formally introduce participation constraints, one must define every agent's outside option, i.e., the payoff of non-participating in the mechanism. The mechanism is then individually rational if no agent suffers a decrease in expected utility relative to this outside option.[24]

A natural way to introduce an outside option might be as follows.[25] An agent can refuse to pay the transfers, but his type and, therefore, the payoff is updated according to the transition probabilities of the setup. The mechanism chooses the policy of the public decision so that it maximizes the expected sum of the participating agents, ignoring the incentives of the non-participating agents. This version of individual rationality implies that, at each round, after the agent observes his type, he can choose whether to participate or not. If he chooses not to participate, then the transfer to him is zero. We extend GUM for this extended model as follows. Let $P_t \subseteq \mathcal{I}$ denote the set of participating agents in round t. The mechanism always follows a decision policy $\chi(P_t, t)$, which maximizes the total expected flow payoff of the participating agents:

$$\chi(P_t, t) \in \arg\max_{\chi'} \quad \mathbb{E}^{\mu(\hat{\theta}_t, \chi')}\Big[\sum_{i \in \mathcal{I}} \sum_{t'=t}^{T} u^i\big(\tilde{\theta}^i_{t'}, \chi'(P_t, t)(\tilde{\theta}_{t'})\big)\Big]. \tag{42}$$

To formally make the mechanism individually rational, we introduce additional transfers for each period an agent quits the mechanism. Let $\tilde{u}^i_{t-}$ and $\tilde{u}^i_t$ denote the anticipated payoff of i starting from round t and assuming that the set of participating agents for the rest of the game will be P_{t-1} and $P_{t-1} \cap P_t$, respectively. Note that $\sum_{i \in P_{t-1} \cap P_t} \mathbb{E}(\tilde{u}^i_t) \geq \sum_{i \in P_{t-1} \cap P_t} \mathbb{E}(\tilde{u}^i_{t-})$ because in the former summation, the public decisions maximize that total anticipated payoff. We introduce additional transfers as follows:

$$y^{\text{exit}\to i}_t = \mathbb{E}(\tilde{u}^i_{t-}) - \mathbb{E}(\tilde{u}^i_t) + \frac{1}{|P_{t-1} \cap P_t|} \sum_{i \in P_{t-1} \cap P_t} \big(\mathbb{E}(\tilde{u}^i_t) - \mathbb{E}(\tilde{u}^i_{t-})\big). \tag{43}$$

[24] We can analogously define coalitional rationality: no set of agents can get a higher total expected utility by some of them using the outside option.

[25] In public goods applications, the appropriate benchmark from which to measure payoff change might not be uniquely defined. For example, if an agent declines to participate in a mechanism, he might be able to avoid paying transfers. However, he might still be affected by the public decision as well as the decisions of the remaining agents (see Green and Laffont (1978) for a discussion). See also Krishna and Perry (1998) for a discussion of type-dependent participation constraints in efficient mechanisms.

The additional transfers ensure that the truthful agents are compensated for the effect of quitting, which guarantees that the truthful agents get the same guaranteed expected utility no matter what the others do.

The extended GUM with additional transfers guarantees the same anticipated payoff for every truthful agent who does not exit the mechanism, and the utility guarantees sum up to the maximum possible total expected utility of all agents. Therefore, the proof of incentive compatibility and collusion resistance applies to this extended setup when exiting is allowed.[26]

6.2 Observing true past types

While we presented our main results assuming that agents observe only public reports regarding other agents' types, our results remain intact if agents observe the true past types (or noisy signals of them) of other agents. For instance, consider a modified model where some agents $\mathcal{O} \subseteq \mathcal{I}$ observe all the true past types. Namely, at period t, every agent $j \in \mathcal{O}$ observes both the reported $\hat{\theta}^i_{t-1}$ as well as the true type θ^i_{t-1} for every i. In this modified model, all our results remain intact, and truth-telling remains an equilibrium. Similarly, we can allow the agents to observe any informative signals regarding the past types.[27] Moreover, within each period t, we can allow an agent whose report is being updated to observe the true types of agents whose reports have already been updated at t. That is, when a report of agent j from an ordered set $\{0, 1, ..., N\}$ is updated, we can allow agent j to observe the true types θ^k_t of all agents $k < j$.

Our mechanism, therefore, extends the results of the efficient budget-balanced mechanism design into the environment where the agents observe signals regarding the past types in addition to the public reports. For instance, the (unbalanced) team mechanism in Athey and Segal (2013) sustains truth-telling as a subgame-perfect equilibrium if all the private types are publicly observed. However, under the Balanced Team Mechanism, the truthful strategy profile is not a subgame-perfect equilibrium in the modified model where agents observe the true past types of other agents.

The proof will be based on a strengthening of CPGUP (22) as follows. Let $\bar{\bar{\mathcal{S}}}^L$ denote

[26]If we want to allow re-entering, then we need to introduce additional transfers compensating the agents in $P_{t-1} \cap P_t$. The rule is the same as before: the re-entering of j changes the expected total payoff of each other agent i, and j pays the negative of this change. Note that re-entering with a type of constant zero payoff throughout the game causes no transfer. Hence, an agent has the same incentives to re-enter and to report the true type instead of that constant zero payoff type.

[27]The proof of Theorem 1 remains intact even if agents observe noisy signals regarding the past types in addition to the public reports.

the set of collusive strategies of $L \subseteq \mathcal{I}$ also allowing them to observe the true past types of others. Namely, a pure extended collusive strategy of a set of agents $L \subseteq \mathcal{I}$ is given by $\bar{\bar{s}}^L = \{\bar{\bar{s}}_t^L\}_{t=1}^T$ where

$$\bar{\bar{s}}_t^L : (\mathbf{\Theta}^{\mathcal{I}})^{t-1} \times \mathbf{\Theta}^L \times \mathcal{H}_{t-1} \to \mathbf{\Theta}^L \tag{44}$$

and the set of the mixed extended collusive strategies are denoted by $\bar{\bar{\mathcal{S}}}^L$. Clearly, $\bar{\mathcal{S}}^L \subset \bar{\bar{\mathcal{S}}}^L$.

In Appendix B we prove that GUM has the following strengthened property, where we replaced $\bar{\mathcal{S}}^{\mathcal{I}\setminus\{i\}}$ with $\bar{\bar{\mathcal{S}}}^{\mathcal{I}\setminus\{i\}}$ in the definition of CPGUP.

$$\sum_{i\in\mathcal{I}} \max_{s^i\in\mathcal{S}^i} \min_{\bar{\bar{s}}^{\mathcal{I}\setminus\{i\}}\in\bar{\bar{\mathcal{S}}}^{\mathcal{I}\setminus\{i\}}} \mathbb{E}\big[U^i(s^i, \bar{\bar{s}}^{\mathcal{I}\setminus\{i\}})\big] = \max_{\bar{\bar{s}}^{\mathcal{I}}\in\bar{\bar{\mathcal{S}}}^{\mathcal{I}}} \sum_{i\in\mathcal{I}} \mathbb{E}\big[U^i(\bar{\bar{s}}^{\mathcal{I}})\big] \tag{45}$$

This proves an even stronger robustness about observing the past of each other: it works for example in any model with noisy and correlated observation, because it only assumes two things.

1. Each agent i can use his strategies as if he did not observe the past of others, meanwhile hiding his current-round type.

2. If i hides his current-round type, then whatever the other agents do must be a legal (mixed) collusive strategy where they fully observe the true past types of all agents.

6.3 Private actions

All our results remain intact if agents take private actions.[28] Namely, we will show that the guaranteed utility property (13) prevails, and as we have seen, all the proven features are consequences of it. In particular, suppose that each period t consists of the following steps:

1. θ_t is chosen from the probability distribution $\prod_{i\in\mathcal{I}\cup\{0\}} \mu(\theta_{t-1}^i, \theta_{t-1}^0, x_{t-1}^i, x_{t-1}^0)$ for every $t > 0$. Each agent $i \in \mathcal{I}$ privately observes θ_t^i, and all agents observe θ_t^0.
2. Each agent makes a report $\hat{\theta}_t^i$.
3. A public decision $x_t^0 \in \boldsymbol{X}$ is made.
4. Each agent $i \in \mathcal{I}$ makes a private decision $x_t^i \in \boldsymbol{X}$.

The notion of a decision policy extends naturally to this extended setup as a vector $\chi' = (\chi'^0, \chi'^1, \ldots, \chi'^N)$ where $\chi'^i \colon \{0, 1, \ldots, T\} \times \mathbf{\Theta} \to \boldsymbol{X}$.[29] A history h and a decision policy χ' determines a probability distribution $\mu(h, \chi')$. The mechanism chooses an efficient

[28]In the working paper version of Athey and Segal (2013), the mechanism includes private decisions.

[29]As in the basic model, the decisions could depend on the full history. Our mechanism will use a decision policy of this simpler form, with no loss in efficiency.

decision policy $\chi = (\chi^0, \chi^1, \dots, \chi^N)$, that is

$$\chi \in \arg\max_{\chi'} \quad \mathbb{E}^{\mu(\theta_0,\chi')}\Big[\sum_{i\in\mathcal{I}}\sum_{t=0}^{T} u^i\big(\tilde{\theta}^i_t, \chi'(\tilde{\theta}_t)\big)\Big]. \tag{46}$$

After all the agents report their types $\hat{\theta}^i_t$, the mechanism makes a recommendation for their private action $\chi^i(\hat{\theta}_t)$ in the future period. When choosing the recommended private actions, the designer calculates the anticipated payoffs. Given the efficient decision policy and the payments described by (10), all the agents report their types truthfully and follow the recommendation when choosing their private actions.[30] This extension of the setup can be combined with Section 6.2. Namely, we can allow the observation of true past types and the observation of past private actions. This yields a similar strengthening of CPGUP (22) as in Section 6.2.

7 Approximate efficiency without transfers

We now consider a modification of our base model without transfers.[31] We will show a transfer-free version of GUM that implements approximate CPGUP with limited error provided that the externalities $\gamma^{i\to j}_t(h)$ defined in (9) are absolutely bounded by a constant D. **Approximate CPGUP with error** R is defined by the following modification of (21).

$$\forall i \in \mathcal{I},\ \forall \bar{\bar{s}}^{\mathcal{I}\setminus\{i\}} \in \bar{\bar{\mathcal{S}}}^{\mathcal{I}\setminus\{i\}}: \qquad \mathbb{E}\big[U^i(s^i_*, \bar{\bar{s}}^{\mathcal{I}\setminus\{i\}})\big] \geqslant C^i; \tag{47}$$

$$\sup_{\bar{\bar{s}}^{\mathcal{I}}\in\bar{\bar{\mathcal{S}}}^{\mathcal{I}}} \sum_{i\in\mathcal{I}} \mathbb{E}\big[U^i(\bar{\bar{s}}^{\mathcal{I}})\big] - \sum_{i\in\mathcal{I}} C^i \leq R \tag{48}$$

We define the **transfer-free GUM** with bound B as the GUM with the following two modifications:

1. There are no transfers, $y \equiv 0$.
2. For any ordered pair of agents $i \neq j$ and for any round t_0, if $\sum_{t=1}^{t_0} \gamma^{i\to j}_t < -B$, then the reports $\hat{\theta}^i_t$ for $t > t_0$ are replaced with random types from $\mu^i\big(\hat{\theta}^i_{t-1}, x_{t-1}\big)$ until the game ends.[32]

In other words, the agents are playing GUM except that the payments are not paid but each ordered pair of agents $(i, j) \in \mathcal{I}^2$ has a separate virtual budget for the ordered bilateral

[30] See Csóka et al. (2023) for a general model with both moral hazard and adverse selection.

[31] The modification is applicable in the environment with private actions and observability of past types.

[32] Note that the replacement of reports with random ones affects both the public decisions x_t and the virtual transfers of $\gamma^{j\to k}_t$ for $j \geq i$ and $t > t_0$.

payments $\gamma_t^{i\to j}$, and if i owes j more than this limit B, then i's reports are replaced with random reports after that period.

The maximum externality a report can cause to another agent is formalized as follows.[33]

$$D := \max_{i,j,t,h_{t,i}} \left|\gamma_t^{i\to j}(h_{t,i})\right| \leq \max_{t,\theta_t,\theta'_t,i} \left(\mathbb{E}^{\mu(\theta_t,\chi)}\Big[\sum_{t'=t}^{T} u^i\big(\tilde{\theta}^i_{t'}, \chi(\tilde{\theta}_{t'})\big)\Big] - \mathbb{E}^{\mu(\theta'_t,\chi)}\Big[\sum_{t'=t}^{T} u^i\big(\tilde{\theta}^i_{t'}, \chi(\tilde{\theta}_{t'})\big)\Big] \right) \tag{49}$$

The inequality tells that D is bounded by the maximum change in the anticipated payoff of any player caused by changing the current type vector.

Theorem 5. *In any game with N agents and T rounds and a bound D as in (49), the transfer-free GUM with $B = D \cdot \sqrt{T \cdot \ln(T \cdot N)}$ implements approximate CPGUP with error bound $R = N^2 \cdot D \cdot \left(1 + \sqrt{T \cdot \ln(T \cdot N)}\right)$.*

We can illustrate the theorem with repeated games where independent copies of the same single-round setup are played in a (large number) T of independent rounds. In this case, the externality is bounded directly by the payoff function as follows.

$$D \leq \sup_{i,\ \theta_i} \left[\sup_x u_i(\theta_i, x) - \inf_x u_i(\theta_i, x)\right]. \tag{50}$$

As a consequence, the per period error bound is at most $N^2 \cdot D \cdot \sqrt{\frac{\ln(T\cdot N)}{T}}$ which converges to 0 as $T \to \infty$.

We note that if the evolution of the type space induced by μ and χ is a finite irreducible Markov chain, then D is a bounded function of T. Therefore, for such setups, the per period error also converges to 0.

We also note that all of these results are valid in the extensions where agents can observe past types and the agents have private actions (see Section 6).

The mechanism and the error bounds could be improved by changing the decision policy as compensation for a "virtual debt". This technique could be particularly efficient in a case of continuous payoff functions and transition rules on connected topological type spaces, in which case the factor $\sqrt{T \ln(T \cdot N)}$ in Theorem 5 could be replaced by $\ln(T)$. We leave this for future research.

The repeated game problem without transfers we considered was solved for constant-size type spaces in Jackson and Sonnenschein (2007) (hereafter, JS).[34] The key idea in JS, is to

[33]Notice that $\mu\big((h_{t-1,N}, \theta_t), \chi\big)$ starting from round t is independent of $h_{t-1,N}$, therefore, we denote this distribution by $\mu(\theta_t, \chi)$.

[34]See also Ball et al. (2022) for a correction of the approximate truthfulness definition used in JS.

use statistical tests to detect lies. If the reported statistics of a coalition of agents differ significantly from their expected true distribution, then the mechanism "punishes" those agents by replacing their reports with random ones. A truthful agent secures a sequence of payoffs that converge to the target level under the efficient decision policy. However, the error term in JS is heavily increasing with the size of type spaces and the number of agents.[35] In particular, the approach used in JS is not applicable if type spaces are infinite. In contrast, our transfer-free GUM achieves efficiency with an error that does not depend on the type space size.

Proof of Theorem 5. For each agent i, consider the joint distribution of the reports of others, as a function of the reports of i. By reports, we mean the random replacement in the cases when the agent already violated the budget constraint. This distribution can be simulated by a collusive strategy $\bar{\bar{s}}^{\mathcal{I}\setminus\{i\}} \in \bar{\bar{\mathcal{S}}}^{\mathcal{I}\setminus\{i\}}$ in GUM. If agent i always reports truthfully, then GUM guarantees $E\big[U^i\big] \geq C^{GUM,i}$ for any collusive strategy of the other agents, including $\bar{\bar{s}}^{\mathcal{I}\setminus\{i\}}$. The two modifications of GUM to the transfer-free GUM change this expected utility as follows.

- In the transfer-free GUM, agent i may suffer a loss by not receiving the transfers. If i is truthful, then $\mathbb{E}_{\theta_t^i}\big[\gamma_t^{i\to j}(\theta_t^i, \hat{\theta}_t^{-i})\big] = 0$ for any fixed report profile $\hat{\theta}_t^{-i}$. About the sequence $\sum\limits_{t=1}^{t_0} \gamma_t^{j\to i}$, once it goes below $-B$, the further terms become 0 in expectation. To sum up, the total expected transfer to i in GUM is at most $(N-1)\cdot(B+D)$.

- The loss due to the obligatory misreporting of i. Hoeffding's inequality shows that $\sum_{t=1}^{t_0} \gamma_t^{i\to j} < -B$ occurs with probability at most $e^{\frac{-2B^2}{t_0\cdot D^2}} \leq e^{\frac{-2\cdot D^2\cdot T\cdot\ln(T\cdot N)}{T\cdot D^2}} = e^{-2\cdot\ln(T\cdot N)} = \frac{1}{T^2\cdot N^2}$. Therefore, the probability that it happens with any of the agents $j \neq i$ at any time point t_0 is less than $\frac{1}{T\cdot N}$. In case it happens, the loss of i due to the obligatory random misreporting is bounded by D per round, which is at most $T\cdot D$ loss in total. Therefore, the expected loss is less than $\frac{D}{N}$.

Therefore, the guaranteed payoff of i is $C^i \geq C^{GUM,i} - \frac{D}{N} - (N-1)\cdot(B+D)$. So the sum of the guaranteed payoffs satisfies the following.

$$\sup_{\bar{\bar{s}}^{\mathcal{I}} \in \bar{\bar{\mathcal{S}}}^{\mathcal{I}}} \sum_{i\in\mathcal{I}} \mathbb{E}\big[U^i(\bar{\bar{s}}^{\mathcal{I}})\big] - \sum_{i\in\mathcal{I}} C^i = \sum_{i\in\mathcal{I}} (C^{GUM,i} - C^i) \leq D + N(N-1)\cdot(B+D) < N^2\cdot(B+D) = R.$$

□

[35] Jackson and Sonnenschein (2007) do not provide explicit bounds for the dependence of the error term on the parameters.

8 Concluding Remarks

We construct an efficient dynamic mechanism with agents privately observing types over time. Transfers are calculated based on the reported types. The reports are updated sequentially within each period. After agent i's report is updated, i pays to j the negative of the externality his report imposes on j. All Bayesian Nash equilibria sustained by our mechanism are utility equivalent. We prove that each agent can guarantee himself an expected utility by being truthful, and those guaranteed utilities add up to the total ex-ante expected surplus. Consequently, our mechanism is collusion-proof, and all Bayesian Nash equilibria are utility equivalent. While we used a standard private values setup with independent types to present our main results, our mechanism remains intact in a more general environment. For instance, the truthful equilibrium sustained by our mechanism remains intact if agents can also choose private actions. In addition, our results remain intact if agents observe signals of the past true types of other agents. Finally, in environments without monetary transfers, we construct a transfer-free GUM that approximately implements CPGUP.

Appendix

A Proof of the additivity of externality

We show equation (31) in Section 5. First, using (9), the sum of the expected externalities across all the agents is equal to:

$$\begin{aligned}\sum_{i\in\mathcal{I}}\gamma_t^{i\to j} &= \sum_{i\in\mathcal{I}}\left(\sum_{t'=t}^{T}\left(\mathbb{E}^{\mu(h_{t,i},\chi)}\left[u^j\left(\tilde{\theta}_{t'}^j,\chi(\tilde{\theta}_{t'})\right)\right]-\mathbb{E}^{\mu(h_{t,i-1},\chi)}\left[u^j\left(\tilde{\theta}_{t'}^j,\chi(\tilde{\theta}_{t'})\right)\right]\right)\right)\\ &= \sum_{t'=t}^{T}\left(\sum_{i\in\mathcal{I}}\left(\mathbb{E}^{\mu(h_{t,i},\chi)}\left[u^j\left(\tilde{\theta}_{t'}^j,\chi(\tilde{\theta}_{t'})\right)\right]-\mathbb{E}^{\mu(h_{t,i-1},\chi)}\left[u^j\left(\tilde{\theta}_{t'}^j,\chi(\tilde{\theta}_{t'})\right)\right]\right)\right).\end{aligned}\tag{51}$$

Second, using the definition of $h_{t,i}$ in (6), the expression inside the $\sum_{t'=t}^{T}()$ on the right-hand side of (51) can be rewritten as:

$$\begin{aligned}&\sum_{i\in\mathcal{I}}\left(\mathbb{E}^{\mu(h_{t,i},\chi)}\left[u^j\left(\tilde{\theta}_{t'}^j,\chi(\tilde{\theta}_{t'})\right)\right]-\mathbb{E}^{\mu(h_{t,i-1},\chi)}\left[u^j\left(\tilde{\theta}_{t'}^j,\chi(\tilde{\theta}_{t'})\right)\right]\right)\\ &=\mathbb{E}^{\mu(h_{t,0},\chi)}\left[u^j\left(\tilde{\theta}_{t'}^j,\chi(\tilde{\theta}_{t'})\right)\right]-\mathbb{E}^{\mu(h_{t,-1},\chi)}\left[u^j\left(\tilde{\theta}_{t'}^j,\chi(\tilde{\theta}_{t'})\right)\right]+\cdots\\ &\quad+\mathbb{E}^{\mu(h_{t,N},\chi)}\left[u^j\left(\tilde{\theta}_{t'}^j,\chi(\tilde{\theta}_{t'})\right)\right]-\mathbb{E}^{\mu(h_{t,N-1},\chi)}\left[u^j\left(\tilde{\theta}_{t'}^j,\chi(\tilde{\theta}_{t'})\right)\right]\\ &=\mathbb{E}^{\mu(h_{t,N},\chi)}\left[u^j\left(\tilde{\theta}_{t'}^j,\chi(\tilde{\theta}_{t'})\right)\right]-\mathbb{E}^{\mu(h_{t,-1},\chi)}\left[u^j\left(\tilde{\theta}_{t'}^j,\chi(\tilde{\theta}_{t'})\right)\right],\end{aligned}\tag{52}$$

where $h_{t,-1}=h_{t-1,N}$.

Finally, using that $h_{t,-1}=h_{t-1,N}$, the expression on the right-hand side of (52) becomes:

$$\begin{aligned}&\mathbb{E}^{\mu(h_{t,N},\chi)}\left[u^j\left(\tilde{\theta}_{t'}^j,\chi(\tilde{\theta}_{t'})\right)\right]-\mathbb{E}^{\mu(h_{t,-1},\chi)}\left[u^j\left(\tilde{\theta}_{t'}^j,\chi(\tilde{\theta}_{t'})\right)\right]\\ &=\mathbb{E}^{\mu(h_{t,N},\chi)}\left[u^j\left(\tilde{\theta}_{t'}^j,\chi(\tilde{\theta}_{t'})\right)\right]-\mathbb{E}^{\mu(h_{t-1,N},\chi)}\left[u^j\left(\tilde{\theta}_{t'}^j,\chi(\tilde{\theta}_{t'})\right)\right].\end{aligned}\tag{53}$$

Thus, we proved that the externality defined in (9) is additive:

$$\sum_{i\in\mathcal{I}}\gamma_t^{i\to j}=\sum_{t'=t}^{T}\left(\mathbb{E}^{\mu(h_{t,N},\chi)}\left[u^j\left(\tilde{\theta}_{t'}^j,\chi(\tilde{\theta}_{t'})\right)\right]-\mathbb{E}^{\mu(h_{t-1,N},\chi)}\left[u^j\left(\tilde{\theta}_{t'}^j,\chi(\tilde{\theta}_{t'})\right)\right]\right)=\gamma_t^{\mathcal{I}\to j}.$$

Q.E.D.

B Proof of Theorem 1

We prove that Theorem 1 holds even in the extended model where the agents may have private decisions (Section 6.3) and may observe the past of each other (Section 6.2). We show that GUM satisfies CPGUP, which is equivalent to collusion-proofness according to Proposition 2.

We prove that every agent $i \in \mathcal{I}$ obtains a utility guarantee of $\Upsilon^i_{0,N}$ by being truthful and $\sum_{i\in\mathcal{I}} \Upsilon^i_{0,N}$ is equal to the highest possible expected total utility.[36] The latter statement is obvious. If all agents are truthful, then by the definition of $\Upsilon^i_{0,N}$, the expected utility of agent i is exactly $\Upsilon^i_{0,N}$. By the definition of the efficient decision policy, if the types are truthfully reported, then the mechanism maximizes the expected total payoff, and the total payoff equals the total utility because the transfers are balanced. Formally,

$$\sum_{i\in\mathcal{I}} \mathbb{E}\Big[U^i\big(\bar{\bar{s}}^{\mathcal{I}}\big)\Big] \leq \sum_{i\in\mathcal{I}} \mathbb{E}^{\mu(\theta_0,\chi)}\Big[\,U^i\big(\tilde{\theta}^i, \chi(\tilde{\theta}), y^i(\tilde{\theta})\big)\,\Big] \tag{54}$$

$$= \sum_{i\in\mathcal{I}} \mathbb{E}^{\mu(\theta_0,\chi)}\left[\sum_{t=0}^{T}\left[u^i(\tilde{\theta}^i_t, \chi(\tilde{\theta}_t)) + y^i_t(\tilde{\theta})\right]\right] = \sum_{i\in\mathcal{I}} \mathbb{E}^{\mu(\theta_0,\chi)}\left[\sum_{t=0}^{T} u^i(\tilde{\theta}^i_t, \chi(\tilde{\theta}_t))\right] = \sum_{i\in\mathcal{I}} \Upsilon^i_{0,N}. \tag{55}$$

Therefore, we only need to prove the following.

$$\mathbb{E}\Big[\,U^i\big(\theta^i, \chi(\theta), y^i(\theta)\big)(s^i, \bar{\bar{s}}^{\mathcal{I}\setminus\{i\}})\,\Big] = \Upsilon^i_{0,N} \qquad \forall i \in \mathcal{I},\ \forall \bar{\bar{s}}^{\mathcal{I}\setminus\{i\}} \in \bar{\bar{\mathcal{S}}}^{\mathcal{I}\setminus\{i\}}. \tag{56}$$

From now on, the strategy s^i is fixed as the truthful strategy of agent i, and a collusive joint strategy $\bar{\bar{s}}^{\mathcal{I}\setminus\{i\}}$ is also fixed. The transition probabilities in the stochastic process depend on the choice of $\bar{\bar{s}}^{\mathcal{I}\setminus\{i\}}$.

Starting from the first round ($t = 1$), we consider the private actions to be made at time $(t, N+1)$. Hence, the Nth player's report of $\hat{\theta}^N_t$ at time (t, N) is followed by the private actions at time $(t, N+1)$, which is followed by the report of $\hat{\theta}^0_{t+1}$ at time $(t+1, 0)$.

[36] Note that $\Upsilon^i_{0,N}$ is only a function of the fixed initial type vector θ_0.

We split each transfer to a sum $y_t^i = \sum_{j=0}^{N+1} y_{t,j}^i$ in the following way.

$$y_{t,j}^i = \begin{cases} 0 & \text{, if } j = 0 \text{ or } j = N+1 \\ \sum_{k \neq i} \gamma_t^{i \to k} & \text{, if } j = i \\ -\gamma_t^{j \to i} & \text{, if } j \neq 0, j \neq N+1 \text{ and } j \neq i \end{cases} \quad . \tag{57}$$

Notice that $y_{t,j}^i$ is a function of $h_{t,j}$.

Let $\mathcal{F}_{t,j}^i$ contain

1. the current history $h_{t,j}$,
2. all public and private decisions up to round $t-1$, namely, $(x_{t'}^i)$ for all $i \in \mathcal{I} \cup \{0\}$ and $t' \leq t-1$,
3. the true current types θ_t, excluding θ_t^i for $j < i$, namely, $\begin{cases} \theta_t^{-i} \text{ if } j < i; \\ \theta_t \text{ if } j \geq i. \end{cases}$

This expresses an upper bound for the total information of all agents excluding i at time (t, j). More formally, we assume that if i does not collude, then $\hat{\theta}_t^j$ must be $\mathcal{F}_{t,j}^i$-measurable.

We prove that the sum of the anticipated payoff $\Upsilon_{t,j}^i$ of the truthful agent i and the total transfers $\sum_{(t',i)=(1,0)}^{(t,j)} y_{t',i}^i$ up to a time (t, j) is a martingale:

$$\mathbb{E}\left[\Upsilon_{t,j}^i + \sum_{(t',i)=(1,0)}^{(t,j)} y_{t',i}^i \,\middle|\, \mathcal{F}_{t,j-1}^i \right] = \Upsilon_{t,j-1}^i + \sum_{(t',i)=(1,0)}^{(t,j-1)} y_{t',i}^i, \tag{58}$$

which is equivalent to:

$$\mathbb{E}\left[\Upsilon_{t,j}^i - \Upsilon_{t,j-1}^i + y_{t,j}^i \,\middle|\, \mathcal{F}_{t,j-1}^i \right] = 0, \tag{59}$$

or alternatively

$$\mathbb{E}\left[\gamma_t^{j \to i} + y_{t,j}^i \,\middle|\, \mathcal{F}_{t,j-1}^i \right] = 0. \tag{60}$$

When the second lower index is $j = 0$, then $(t, j-1) = (t, -1)$ is identified by the time $(t-1, N+1)$. We now formally prove that condition (60) holds. There are four cases depending on the type of report being updated.

Case 1 ($j = 0$, updating the public report). Since $y_{t,0}^i = 0$, we need to show that

$\mathbb{E}\left[\gamma_t^{0\to i} \,\middle|\, \mathcal{F}_{t,-1}^i\right] = 0$; see (60). That is, by the definition of externality (9), we need

$$\mathbb{E}\left[\mathbb{E}^{\mu(h_{t,0},\chi)}\left[\sum_{t'} u^i\big(\tilde{\theta}_{t'}^i, \chi(\tilde{\theta}_{t'})\big)\right] \,\middle|\, \mathcal{F}_{t,-1}^i\right] = \mathbb{E}\left[\mathbb{E}^{\mu(h_{t,-1},\chi)}\left[\sum_{t'} u^i\big(\tilde{\theta}_{t'}^i, \chi(\tilde{\theta}_{t'})\big)\right] \,\middle|\, \mathcal{F}_{t,-1}^i\right]. \quad (61)$$

The left-hand side of the equality is the total expectation of $\sum_{t'} u^i\big(\tilde{\theta}_{t'}^i, \chi(\tilde{\theta}_{t'})\big)$ as θ_t ranges through $\boldsymbol{\Theta}^{N+1}$ according to the transition probabilities governed by μ. Hence, (61) follows from the law of total expectation:

$$\mathbb{E}\left[\mathbb{E}^{\tilde{\theta}\sim\mu(h_{t,0},\chi)}\left[\sum_{t'} u^i\big(\tilde{\theta}_{t'}^i, \chi(\tilde{\theta}_{t'})\big)\right] \,\middle|\, \mathcal{F}_{t,-1}^i\right] \quad (62)$$

$$= \mathbb{E}\left[\mathbb{E}^{\hat{\theta}_t^0\sim\mu^0(\hat{\theta}_{t-1}^0, x_{t-1})}\left[\mathbb{E}^{\tilde{\theta}\sim\mu((h_{t,-1},\hat{\theta}_t^0),\chi)}\left[\sum_{t'} u^i\big(\tilde{\theta}_{t'}^i, \chi(\tilde{\theta}_{t'})\big)\right] \,\middle|\, \mathcal{F}_{t,-1}^i\right] \,\middle|\, \mathcal{F}_{t,-1}^i\right] \quad (63)$$

$$= \mathbb{E}\left[\mathbb{E}^{\tilde{\theta}\sim\mu(h_{t,-1},\chi)}\left[\sum_{t'} u^i\big(\tilde{\theta}_{t'}^i, \chi(\tilde{\theta}_{t'})\big)\right] \,\middle|\, \mathcal{F}_{t,-1}^i\right]. \quad (64)$$

Case 2 ($j = i$, updating agent i's report). Since $j = i$, we have $\gamma_t^{j\to i} + y_{t,j}^i = \gamma_t^{i\to i} + \sum_{k\neq i} \gamma_t^{i\to k} = \sum_k \gamma_t^{i\to k}$. Thus according to (60), we need to show that $\mathbb{E}\left[\sum_k \gamma_t^{i\to k} \,\middle|\, \mathcal{F}_{t,i-1}^i\right] = 0$. That is, we need

$$\sum_k \mathbb{E}\left[\mathbb{E}^{\mu(h_{t,i},\chi)}\left[\sum_{t'} u^k\big(\tilde{\theta}_{t'}^k, \chi(\tilde{\theta}_{t'})\big)\right] - \mathbb{E}^{\mu(h_{t,i-1},\chi)}\left[\sum_{t'} u^k\big(\tilde{\theta}_{t'}^k, \chi(\tilde{\theta}_{t'})\big)\right] \,\middle|\, \mathcal{F}_{t,i-1}^i\right] = 0.$$

Hence, it is enough to show that for each $k \in \mathcal{I}$ we have

$$\mathbb{E}\left[\mathbb{E}^{\mu(h_{t,i},\chi)}\left[\sum_{t'} u^k\big(\tilde{\theta}_{t'}^k, \chi(\tilde{\theta}_{t'})\big)\right] \,\middle|\, \mathcal{F}_{t,i-1}^i\right] = \mathbb{E}\left[\mathbb{E}^{\mu(h_{t,i-1},\chi)}\left[\sum_{t'} u^k\big(\tilde{\theta}_{t'}^k, \chi(\tilde{\theta}_{t'})\big)\right] \,\middle|\, \mathcal{F}_{t,i-1}^i\right]. \quad (65)$$

As agent i is truthful, his report is updated from his past true type $\hat{\theta}_{t-1}^i = \theta_{t-1}^i$ to his new true type $\hat{\theta}_t^i = \theta_t^i$. Note that because the true type θ_t^i is revealed at time (t, i), we can view $\hat{\theta}_t^i$ as a random variable governed by μ. By the law of iterated expectations, we have

$$\mathbb{E}\left[\mathbb{E}^{\tilde{\theta}\sim\mu(h_{t,i},\chi)}\left[\sum_{t'} u^k\big(\tilde{\theta}_{t'}^k, \chi(\tilde{\theta}_{t'})\big)\right] \,\middle|\, \mathcal{F}_{t,i-1}^i\right] \quad (66)$$

$$= \mathbb{E}\left[\mathbb{E}^{\hat{\theta}_t^i\sim\mu^i(\hat{\theta}_{t-1}^0, \hat{\theta}_{t-1}^i, x_{t-1})}\left[\mathbb{E}^{\tilde{\theta}\sim\mu((h_{t,i-1},\hat{\theta}_t^i),\chi)}\left[\sum_{t'} u^k\big(\tilde{\theta}_{t'}^k, \chi(\tilde{\theta}_{t'})\big)\right] \,\middle|\, \mathcal{F}_{t,i-1}^i\right] \,\middle|\, \mathcal{F}_{t,i-1}^i\right] \quad (67)$$

$$= \mathbb{E}\Bigg[\mathbb{E}^{\tilde{\theta} \sim \mu(h_{t,i-1},\chi)} \Bigg[\sum_{t'} u^k\big(\tilde{\theta}^k_{t'}, \chi(\tilde{\theta}_{t'})\big) \Bigg] \Bigg| \mathcal{F}^i_{t,i-1} \Bigg]. \tag{68}$$

Case 3 ($j \neq i$ and $j \neq 0$, updating other agent j's report). By (57), we have $\gamma_t^{j \to i} + y^i_{t,j} = 0$, and thus (60) holds trivially.

Case 4 ($j = N+1$, making the private actions). Since $y^i_{t,N+1} = 0$, we need to show that $\mathbb{E}\Big[\gamma_t^{N+1 \to i} \Big| \mathcal{F}^i_{t,-1} \Big] = 0$; see (60). That is, by the definition of externality (9), we need

$$\mathbb{E}\Bigg[\mathbb{E}^{\mu(h_{t,N+1},\chi)} \Bigg[\sum_{t'} u^i\big(\tilde{\theta}^i_{t'}, \chi(\tilde{\theta}_{t'})\big) \Bigg] \Bigg| \mathcal{F}^i_{t,N} \Bigg] = \mathbb{E}\Bigg[\mathbb{E}^{\mu(h_{t,N},\chi)} \Bigg[\sum_{t'} u^i\big(\tilde{\theta}^i_{t'}, \chi(\tilde{\theta}_{t'})\big) \Bigg] \Bigg| \mathcal{F}^i_{t,N} \Bigg]. \tag{69}$$

The distributions $\mu(h_{t,N+1}, \chi)$ and $\mu(h_{t,N}, \chi)$ are the same, since at time (t, N) only deterministic presumed private actions are made according to the decision policy χ.

Now as we proved that $\Upsilon^i_{t,j} + \sum_{(t',j')=(1,0)}^{(t,j)} y^i_{t',j'}$ is a martingale, we conclude that every agent can obtain a utility guarantee of $\Upsilon^i_{0,N}$ by being truthful, $\Upsilon^i_{0,N} = \mathbb{E}^\mu\Big[\Upsilon^i_{0,N} + \sum_{(t,j)=(1,0)}^{(0,N)} y^i_{t,j}\Big]$

$$= \mathbb{E}\Bigg[\Upsilon^i_{T,N} + \sum_{(t,j)=(1,0)}^{(T,N)} y^i_{t,j}\Bigg] = \mathbb{E}\Bigg[\sum_{k=0}^{T} u^i\big(\tilde{\theta}^i_k, \chi(\tilde{\theta}_k)\big) + \sum_{t=0}^{N} y^i_t\Bigg] = \mathbb{E}[U^i] = \mathbb{E}\big[U^i(s^i, \bar{\bar{s}}^{\mathcal{I}\setminus\{i\}})\big]. \tag{70}$$

Q.E.D.

C Weaknesses of Che–Kim's Strong Collusion-Proofness and Safronov's Coalition-proofness

Our definition of collusion-proofness implies Che–Kim's and Safronov's definitions as well as all earlier definitions of collusion-proofness we are aware of. Equivalently, a violation of any of these definitions of collusion-proofness is a weak collusive equilibrium which changes the expected utility of at least one player. In this appendix, we will elaborate on the other direction: Che–Kim's and Safronov's definitions have behaviourally relevant weaknesses compared to our definition.

All definitions of collusion-proofness we are aware of were defined for static games. The simplest setup about collusion has two agents and a principal, considering the possible coalition of the two agents. For this case, Laffont and Martimort (2000) (following Cramton and Palfrey (1995)) introduced a definition of "weak collusion-proofness". Then Che and

Kim used the analogous definition for multiple agents, called "strong collusion proofness". Hereby, this definition did not address some additional potential problems with collusion, as we will shortly demonstrate. These definitions did not allow the agents to bond themselves about what they would report.[37] They also did not consider the possibility that the agents may reveal and verify their true types. Che–Kim's strong collusion-proofness was entirely focusing on the utility of the principal. In particular, in an arbitrary game with a principal and multiple agents, if the payoff of the principal is a constant, then every Nash equilibrium is "Che–Kim strongly collusion-proof". However, if the principal also gets a higher payoff by a rational collusion of the agents, then it was a violation of Che–Kim's strong collusion-proofness. They used it for a game with a distinguished principal and at some agents, where two agents were distinguished. They considered only those coalitions that included those two agents and did not include the principal.

In a setup similar to ours but in a static framework, Safronov (2018) introduced the simpler definition of coalition-proofness meaning that no set of agents can benefit in total by a collusive deviation. This definition has no distinguished agent, and this is a reason why there is no implication between Safronov's coalition-proofness and Che–Kim's strong collusion-proofness.

Safronov's Coalition-proofness (Safronov, 2018). In a static game G with a type profile θ, a strategy profile s_* is Safronov's coalition-proof if for any type profile θ, any coalition $C \subset \mathcal{I}$ and any joint strategy s^C of agents in C, the following condition holds:

$$\sum_{i\in C} \mathbb{E}\big[(U^i(s^C, s_*^{\mathcal{I}\setminus C})\big] \leqslant \sum_{i\in C} \mathbb{E}\big[U^i(s_*^{\mathcal{I}})\big]. \tag{71}$$

Safronov's coalition-proofness can be interpreted as an extension of Nash equilibrium to groups of colluding agents. We could justify this definition so that if the coalition can benefit from a deviation, then an appropriate collusive transfer rule can make each agent benefit from colluding.

While Safronov's mechanism – as a special case of the GUM – satisfies our stronger definition of collusion-proofness, Safronov's coalition-proofness is a weaker term. For instance, it does not imply uniqueness. For example, there are two coalition-proof equilibria in the symmetrical battle of the sexes. Safronov's coalition-proofness does not automatically extend to a dynamic environment where agents learn private information after some actions.

[37] Che–Kim had an illustrative example of a knockout auction that bonds them how to behave later, but it did not meet their own definition of collusion. They added a footnote that the specific knockout auction in that specific example has an equivalent outcome as another behavior that meets their definition of collusion.

Moreover, Safronov's coalition-proofness is susceptible to the agents' beliefs regarding the formed coalitions.

We next provide an example that satisfies Safronov's coalition-proofness but where agents find it beneficial to collude and coordinate their reports for some beliefs regarding the formed coalitions.

Consider the following three-period game ($T = 3$) with four players, $\mathcal{I} = \{1, 2, 3, 4\}$. The timing is as follows.

$\boldsymbol{t} = \mathbf{1}$: Agent 1 publicly chooses $x_1 \in \{Quit, Bet\}$. If $x_1 = Quit$, then the game ends. If $x_1 = Bet$, then the game proceeds to the next period.

$\boldsymbol{t} = \mathbf{2}$: Agent 2 is trying to predict what Agent 3 and Agent 4 will do in a prisoners' dilemma they play at $t = 3$, described below. Formally, Agent 2 privately chooses a "guess" $x_2 = (g_3, g_4) \in \{Y, N\} \times \{Y, N\}$.

$\boldsymbol{t} = \mathbf{3}$: Agent 3 and Agent 4 play a prisoners' dilemma, i.e., make simultaneous choices: Agent 3 chooses $x_3 \in \{Y, N\}$ and Agent 4 chooses $x_4 \in \{Y, N\}$.

Utilities. The utilities are as follows. If $x_1 = Quit$, then all the players collect zero utilities $(0, 0, 0, 0)$. If $x_1 = Bet$, then the utilities are determined by actions of Agent 2, Agent 3, and Agent 4. We distinguish two scenarios. First, if Agent 2's guess is *correct* $\big(g_3 = x_3$ and $g_4 = x_4\big)$. The utilities are summarized in Table 4. Second, if Agent 2's guess is *wrong* $\big($either $g_3 \neq x_3$, or $g_4 \neq x_4$, or both$\big)$. The utilities are summarized in Table 5. Intuitively, Agent 3 and Agent 4 play the prisoner's dilemma (N is analogous to "cooperate" and Y is analogous to "defect"). Finally, Agent 2 is better off if his guess is correct, whereas Agent 1 is better off if Agent 2's guess is wrong.

u_1, u_2, u_3, u_4	$x_4 = Y$	$x_4 = N$
$x_3 = Y$	$-2, 1, -10, -10$	$-2, 1, 10, -20$
$x_3 = N$	$-2, 1, -20, 10$	$-2, 1, 0, 0$

Table 3: utilities if Agent 2's guess is *correct* $\big(g_3 = x_3$ and $g_4 = x_4\big)$.

u_1, u_2, u_3, u_4	$x_4 = Y$	$x_4 = N$
$x_3 = Y$	$8, -9, -10, -10$	$8, -9, 10, -20$
$x_3 = N$	$8, -9, -20, 10$	$8, -9, 0, 0$

Table 4: utilities if Agent 2's guess is *wrong* $\big($either $g_3 \neq x_3$ or $g_4 \neq x_4$, or both$\big)$.

Equilibrium analysis (without collusion). To characterize the equilibrium, note that there is a unique equilibrium (Y, Y) in the Agent 3 and Agent 4 prisoner's dilemma at $t = 3$. At

$t = 2$, Agent 2 correctly anticipates the outcome at $t = 3$ and chooses a guess $x_2 = (Y, Y)$. At $t = 1$, Agent 1 optimally chooses $x_1 = Quit$ anticipating a negative utility from $x_1 = Bet$. Formally, the unique subgame perfect Nash equilibrium is[38]

$$x_1 = Quit,\ x_2 = (Y, Y),\ x_3 = Y,\ \text{and}\ x_4 = Y. \tag{72}$$

Collusion analysis. Let us now consider the possibility of collusion between Agent 3 and Agent 4. Since they are playing a prisoner's dilemma, they have incentives to collude and choose $x_3 = N$ and $x_4 = N$. Anticipating the collusion, Agent 2 chooses a guess $x_2 = (N, N)$. If Agent 1 chooses $x_1 = Bet$, he collects -2 if Agent 2's guess is correct and 8 if Agent 2's guess is wrong. Therefore, Agent 1 chooses $x_1 = Bet$ if he believes that Agent 2's guess is wrong with probability at least 0.2. However, by checking all the joint deviations for all possible coalitions, it is straightforward that the equilibrium in (72) is Safronov's coalition-proof.

The example above illustrates that Safronov's coalition-proofness is susceptible to uncertainty regarding formed coalitions. We note that our analysis remains intact if the game was static, i.e., played within the same period. Moreover, one could construct an example where even if it is known that the agents are open to collusion, it remains uncertain which coalition would be formed.

For example, we could replace the 2-player prisoner's dilemma at $t = 3$ with a 3-player symmetric game between agents 3, 4, and 5 where only two out of the three players should cooperate. Namely, each of agents 3, 4, and 5, chooses either Y or N. Then the utilities from the agents' choices are as follows (symmetrically between the agents): $(-10, -10, -10)$ after (Y, Y, Y); $(-40, 0, 0)$ after (N, Y, Y); $(10, 10, -20)$ after (N, N, Y); and $(-50, -50, -50)$ after (N, N, N). This modification of the initial game has the unique Nash equilibrium (Y, Y, Y) and three symmetrical options for plausible collusion. Therefore, even if Agent 2 knows that agents 3, 4, and 5 are open to collusion, he cannot predict what they will play.

These examples demonstrated that a strategy profile may be the unique Nash equilibrium, this equilibrium is Safronov's coalition-proof, while it is not a rational strategy for an agent due to some uncertainty about the collusion of others. The same issue could be demonstrated also for Che–Kim's strong collusion-proofness by introducing a principal who gets lower utility in the case of the collusive outcome than for the non-collusive outcome.

[38] Note that the same argument proves that $x_1 = Quit$ with probability one in every Nash equilibrium.

D Proof of Theorem 4

In this section, we will specify and slightly modify Example 5.1. Our analysis will be based on (35). In order to get convenient values, we replace the utilities of (Blue, Red, Green) to (84 or 104, 84 or 104, −204), and hereby in (35) the constant factor 150 − 100 = 50 is replaced by 204 − 104 = 100. We use the notation $\delta_t^i = \widehat{p}_t^i - p_t^i$. As this section is dedicated to the BTM, we use the simplified notation γ_t^i for $\gamma_t^{BTM,i}$.

Lemma 2. *Assume that $\widehat{p}_t^{blue} = p_t^{blue}$ in every even round t, and $\widehat{p}_t^{red} = p_t^{red}$ in every odd round t, and both are true in the first and last rounds $t = 0$ and $t = T$. Then*

$$\mathbb{E}(\gamma^{blue}) = \mathbb{E}\Big(\sum_{t=1}^{T} \gamma_t^{blue}\Big) = 100 \cdot \sum_{t=1}^{T} \mathbb{E}\big(-\delta_{t-1}^{red} \cdot \delta_t^{blue}\big) = 100 \cdot \sum_{t=1}^{\lfloor \frac{T-2}{2} \rfloor} \mathbb{E}\big(-\delta_{2t}^{red} \cdot \delta_{2t+1}^{blue}\big) \tag{73}$$

$$\mathbb{E}(\gamma^{red}) = \mathbb{E}\Big(\sum_{t=1}^{T} \gamma_t^{red}\Big) = 100 \cdot \sum_{t=1}^{T} \mathbb{E}\big(-\delta_{t-1}^{blue} \cdot \delta_t^{red}\big) = 100 \cdot \sum_{t=1}^{\lfloor \frac{T-1}{2} \rfloor} \mathbb{E}\big(-\delta_{2t-1}^{blue} \cdot \delta_{2t}^{red}\big) \tag{74}$$

D.1 The setup

Consider now Example 5.1 with $T = 4$ number of rounds (and $N \geq 3$ players). We define a finite space of types for both players. The types are encoded as follows:

[b/r: Blue or Red][Round number]:[the percentage that his type will be HIGH]%.

Types of agent Blue:		*Types of agent Red:*	
b0:50%		r0:50%	
b1:30%,	b1:70%	r1:50%	
b2:20%,	**b2:80%**	r2:20%,	r2:80%
b3:10%,	b3:90%	**r3:10%,**	**r3:90%**
b4:0%,	**b4:100%**	**r4:0%,**	**r4:100%**
(LOW)	(HIGH)	(LOW)	(HIGH)

The transition probabilities can be calculated from (the martingale property of) the percentages. E.g., b1:30% is transitioning to b2:20% or to b2:80% with probabilities 5/6 and 1/6, respectively, because $30\% = \frac{5}{6} \cdot 20\% + \frac{1}{6} \cdot 80\%$.

We make some modifications to the setup in order to incentivize the agents (under the balanced Team Mechanism) to tell the truth in specific rounds (marked with **bold**). Hereby, we will be able to focus on the possibilities of deviations only in the rest of the rounds, which will simplify the analysis.

We add 4+1 extra public decisions, each of them affecting only one agent. In round 2, there is a public decision "b2:20%" or "b2:80%". If this decision does not coincide with the true type of Blue, then Blue gets a payoff -10^{42} (instead of 0) for that round. We do the analogous modification in rounds 2 and 4 for Blue and in rounds 3 and 4 for Red (marked with **bold**). We add one more public decision, for Red in round 2, but this time with an ε loss in payoff.

Formally, we have

$$
\begin{aligned}
X_2 &= \{\text{“b2:20\%”}, \text{“b2:80\%”}\} \times \{\text{“r2:20\%”}, \text{“r2:80\%”}\} \\
X_3 &= \{\text{“r3:10\%”}, \text{“r3:90\%”}\} \\
X_4 &= \{\text{“b4:0\%”}, \text{“b4:100\%”}\} \times \{\text{“r4:0\%”}, \text{“r4:100\%”}\} \times \{\text{“YES”}, \text{“NO”}\} \\
u_1^{blue} = u_1^{red} &\equiv 0 \\
u_2^{blue}(\theta_2^{blue}, x_2) &= -10^{42} \cdot \mathrm{I}(\theta_2^{\mathrm{blue}} \neq \mathrm{x}_{2,1}) \\
u_2^{red}(\theta_2^{red}, x_2) &= -\varepsilon \cdot \mathrm{I}(\theta_2^{\mathrm{red}} \neq \mathrm{x}_{2,2}) \\
u_3^{blue} &\equiv 0 \\
u_3^{red}(\theta_3^{red}, x_3) &= -10^{42} \cdot \mathrm{I}(\theta_3^{\mathrm{red}} \neq \mathrm{x}_{3,1}) \\
u_4^{blue}(\theta_4^{blue}, x_4) &= -10^{42} \cdot \mathrm{I}(\theta_4^{\mathrm{blue}} \neq \mathrm{x}_{4,1}) + \mathrm{I}(\mathrm{x}_{4,3} = \mathrm{YES}) \cdot \big(84 + 20 \cdot \mathrm{I}(\theta_4^{\mathrm{blue}} = \mathrm{HIGH})\big) \\
u_4^{red}(\theta_4^{red}, x_4) &= -10^{42} \cdot \mathrm{I}(\theta_4^{\mathrm{red}} \neq \mathrm{x}_{4,1}) + \mathrm{I}(\mathrm{x}_{4,3} = \mathrm{YES}) \cdot \big(84 + 20 \cdot \mathrm{I}(\theta_4^{\mathrm{red}} = \mathrm{HIGH})\big) \\
u_4^{green}(\theta_4^{green}, x_4) &= -204 \cdot \mathrm{I}(\mathrm{x}_{4,3} = \mathrm{YES})
\end{aligned}
$$

D.2 The analysis

Under the balanced Team Mechanism, these 4+1 extra public decisions induce no transfers between the agents. The only effect is that the agents get a punishment if they do not report truthfully in the specified rounds.

Elimination step 1. Therefore, by weak domination (or a dynamic version of strict domination), we eliminate the possibilities of not reporting the true types at these 4 extra public decisions (marked with **bold**) where the punishments are huge. It leaves a total of 3 binary decisions for the two players: $\widehat{p}_1^{blue}$, $\widehat{p}_2^{red}$ and $\widehat{p}_3^{blue}$. The agents only observe their own types and these earlier decisions of the other agent.

In this reduced game, the utilities $u(x_t, \theta_t^i)$ are unaffected by these three decisions, and the conditions of Lemma 2 apply. Therefore,

- Blue is maximizing $\mathbb{E}\big(\gamma^{blue}\big) - \frac{1}{N-1}\mathbb{E}\big(\gamma^{red}\big) = 100 \cdot \mathbb{E}\big(-\delta_2^{red} \cdot \delta_3^{blue}\big) - \frac{100}{N-1}\mathbb{E}\big(-\delta_1^{blue} \cdot \delta_2^{red}\big)$,
- Red is maximizing $\mathbb{E}\big(\gamma^{red}\big) - \frac{1}{N-1}\mathbb{E}\big(\gamma^{blue}\big) = 100 \cdot \mathbb{E}\big(-\delta_1^{blue} \cdot \delta_2^{red}\big) - \frac{100}{N-1}\mathbb{E}\big(-\delta_2^{red} \cdot \delta_3^{blue}\big)$.

We refer to the binary options and types by "low" and "high", and in this sense, we can say that two decisions are of the "same kind", denoted by "$\sim$", or "opposite", denoted by "$\nsim$". Notice that the reduced game is symmetric in "low" and "high".

Elimination step 2. Let us start with the last decision $\widehat{p}_3^{blue}$. As (73) and (74) show, it can only affect $\mathbb{E}\big(\gamma^{blue}\big) = 100 \cdot \mathbb{E}(-\delta_2^{red} \cdot \delta_3^{blue})$. If $\widehat{p}_2^{red} = 20\%$, then $-\delta_2^{red} \geq 0$, therefore, choosing $\widehat{p}_3^{blue} = 90\%$ weakly dominates $\widehat{p}_3^{blue} = 10\%$. Analogously, if $\widehat{p}_2^{red} = 80\%$, then $-\delta_2^{red} \leq 0$, therefore, choosing $\widehat{p}_3^{blue} = 10\%$ weakly dominates $\widehat{p}_3^{blue} = 90\%$. Hereby we could conclude by weak dominance that Blue should choose $\widehat{p}_3^{blue} \nsim \widehat{p}_2^{red}$.

From now on, the formula γ^{blue} will be calculated as a function of p_1^{blue}, $\widehat{p}_1^{blue}$ and $\widehat{p}_2^{red}$ and with the assumption that $\widehat{p}_3^{blue} \nsim \widehat{p}_2^{red}$. It leaves a total of 2 binary decisions for the two players: $\widehat{p}_1^{blue}$ and $\widehat{p}_2^{red}$. Consider now the second decision $\widehat{p}_2^{red}$. It has an effect on $\mathbb{E}\big(\gamma^{red}\big) = 100 \cdot \mathbb{E}(-\delta_1^{blue} \cdot \delta_2^{red})$ and $\mathbb{E}\big(\gamma^{blue}\big) = 100 \cdot \mathbb{E}(-\delta_2^{red} \cdot \delta_3^{blue})$ and $u_2^{red}(\theta_2^{red}, x_2) = -\varepsilon \cdot \mathrm{I}(\theta_2^{\mathrm{red}} \neq \hat{\theta}_2^{\mathrm{red}})$.

- The effect on $\mathbb{E}\big(\gamma^{red}\big)$. If $\widehat{p}_1^{blue} \nsim p_1^{blue}$, then choosing $\widehat{p}_2^{red} \nsim \widehat{p}_1^{blue}$ increases $\mathbb{E}\big(\gamma^{red}\big)$ by $100 \cdot (0.7 - 0.3) \cdot (0.8 - 0.2) = 24$. If $\widehat{p}_1^{blue} = p_1^{blue}$, then $\mathbb{E}\big(\gamma^{red}\big) = 0$ independently of $\widehat{p}_2^{red}$.
- The effect on $\mathbb{E}\big(\gamma^{blue}\big)$. If $p_3^{blue} \sim p_2^{red}$, then $\widehat{p}_2^{red}$ has no effect on $\mathbb{E}\big(\gamma^{blue}\big)$. But if $p_3^{blue} \nsim p_2^{red}$, then $\widehat{p}_2^{red} \nsim p_2^{red}$ increases $\mathbb{E}\big(\gamma^{blue}\big)$ by $100 \cdot (0.8 - 0.2) \cdot (0.9 - 0.1) = 48$.
- The effect on u_2^{red}: A loss of ε for the case $\widehat{p}_2^{red} \nsim p_2^{red}$.

Thus Red should choose $\widehat{p}_2^{red} \nsim \widehat{p}_1^{blue}$ or $\widehat{p}_2^{red} = p_2^{red}$. Namely, if both decisions would be the same, then Red is strictly better by choosing it. Otherwise, the best choice depends on his belief about the probability that $\widehat{p}_1^{blue} \nsim p_1^{blue}$.

Now the Reader can jump to the matrix games, but we explain in short the dilemma about the first decision $\widehat{p}_1^{blue}$. It can affect both $\mathbb{E}(\gamma^{red})$ and $\mathbb{E}(\gamma^{blue})$. If Red chooses the strategy $\widehat{p}_2^{red} = p_2^{red}$, then it does not matter what Blue does. So consider the case when Red uses his other strategy $\widehat{p}_2^{red} \nsim \widehat{p}_1^{blue}$. In this case, $\mathbb{E}(\gamma^{blue})$ is 48 or 0, and the probabilities depend on $\widehat{p}_1^{blue}$, and also $\Pr\big(\delta_2^{red} \neq 0\big) = 1/2$ independently of $\widehat{p}_1^{blue}$. But for example, if $p_1^{blue} = 30\%$, then they will choose $\widehat{p}_1^{blue} = 70\%$, then $\widehat{p}_2^{red} = 20\%$, and then $\widehat{p}_3^{blue} = 90\%$. Therefore,

$$\Pr\big(\gamma^{blue} = 48\big) = \Pr\big(\delta_3^{blue} \neq 0\big) = \Pr\big(p_3^{blue} = 10\% \;\big|\; p_1^{blue} = 30\%\big) = \frac{90\% - 30\%}{90\% - 10\%} = \frac{3}{4}.$$

With $\widehat{p}_1^{blue} = p_1^{blue} = 30\%$, it would be $P\big(\gamma^{blue} > 0\big) = \Pr\big(p_3^{blue} = 90\% \mid p_1^{blue} = 30\%\big) = 1/4$. This shows that Blue should report $\widehat{p}_1^{blue} = 1 - p_1^{blue}$, and therefore, Red should report the opposite.

Assuming that Blue and Red use a symmetric strategy for high and low, we get the following 2×2 matrix game which only includes γ. Keep in mind that the non-constant terms in the utility of Blue is $\gamma^{blue} - \frac{1}{N-1}\gamma^{red}$ and for Red it is $\gamma^{red} - \frac{1}{N-1}\gamma^{blue} - \varepsilon \cdot \mathrm{I}(\widehat{\mathrm{p}}_2^{\mathrm{red}} \nsim \mathrm{p}_2^{\mathrm{red}})$. (We use a normalization factor of $1/3$ for $\big(\mathbb{E}(\gamma^{blue}),\ \mathbb{E}(\gamma^{red})\big)$ in order to have smaller integers.)

$\frac{1}{3}\big(\mathbb{E}(\gamma^{blue}),\ \mathbb{E}(\gamma^{red})\big)$	$\widehat{p}_2^{red} = p_2^{red}$	$\widehat{p}_2^{red} \nsim \widehat{p}_1^{blue}$
$\widehat{p}_1^{blue} = p_1^{blue}$	$(0,\ 0)$	$(2,\ 0)$
$\widehat{p}_1^{blue} = 1 - p_1^{blue}$	$(0,\ 0)$	$(6,\ 4)$

We can see that Blue prefers $\widehat{p}_1^{blue} = 1 - p_1^{blue}$, and therefore, Red should choose $\widehat{p}_2^{red} \nsim \widehat{p}_1^{blue}$ (if $\frac{\varepsilon}{6} < 4 - \frac{6}{N-1}$).

We could find some arguments to exclude the rationality of asymmetric strategies with respect to high and low, but it is easier to extend the matrix with all asymmetric strategies. The further pure strategies of Blue are always reporting $p_1^{blue} = 30\%$ and always reporting $p_1^{blue} = 70\%$. As for Red, we only need to consider the options when $\widehat{p}_1^{blue} \sim p_2^{red}$ because otherwise he should choose $\widehat{p}_2^{red} = p_2^{red}$. Therefore, the two extra pure strategies of red are the followings.

- "preferably high", meaning that $\widehat{p}_2^{red} = 80\%$ unless if $\big(\widehat{p}_1^{blue} = 70\%$ and $p_2^{red} = 20\%\big)$;
- "preferably low", meaning that $\widehat{p}_2^{red} = 20\%$ unless if $\big(\widehat{p}_1^{blue} = 30\%$ and $p_2^{red} = 80\%\big)$.

Now we get the following 4×4 game.

$\frac{1}{3}\big(\mathbb{E}(\gamma^{blue}),\ \mathbb{E}(\gamma^{red})\big)$	$\widehat{p}_2^{red} = p_2^{red}$	$\widehat{p}_2^{red} \nsim \widehat{p}_1^{blue}$	$\widehat{p}_2^{red}$ pref. high	$\widehat{p}_2^{red}$ pref. low
$\widehat{p}_1^{blue} = p_1^{blue}$	$(0,\ 0)$	$(2,\ 0)$	$(1,\ 0)$	$(1,\ 0)$
$\widehat{p}_1^{blue} = 1 - p_1^{blue}$	$(0,\ 0)$	$(6,\ 4)$	$(3,\ 2)$	$(3,\ 2)$
$\widehat{p}_1^{blue} = 70\%$	$(0,\ 0)$	$(4,\ 2)$	$(1,\ 0)$	$(3,\ 2)$
$\widehat{p}_1^{blue} = 30\%$	$(0,\ 0)$	$(4,\ 2)$	$(3,\ 2)$	$(1,\ 0)$

Or directly the non-constant terms in the utilities for $N = 3$ and $\varepsilon = 3$ are the following, with the normalization factor changed to 4/3.

$\frac{4}{3}\big(\mathbb{E}(U^{blue}),\ \mathbb{E}(U^{red})\big) + C$	$\widehat{p}_2^{red} = p_2^{red}$	$\widehat{p}_2^{red} \not\approx \widehat{p}_1^{blue}$	$\widehat{p}_2^{red}$ pref. high	$\widehat{p}_2^{red}$ pref. low
$\widehat{p}_1^{blue} = p_1^{blue}$	$(0,\ 0)$	$(8,\ -2)$	$(4,\ -1)$	$(4,\ -1)$
$\widehat{p}_1^{blue} = 1 - p_1^{blue}$	$(0,\ 0)$	$(16,\ 2)$	$(8,\ 1)$	$(8,\ 1)$
$\widehat{p}_1^{blue} = 70\%$	$(0,\ 0)$	$(12,\ -2)$	$(4,\ -1)$	$(8,\ 1)$
$\widehat{p}_1^{blue} = 30\%$	$(0,\ 0)$	$(12,\ -2)$	$(8,\ 1)$	$(4,\ -1)$

Elimination step 3. $\widehat{p}_1^{blue} = 1 - p_1^{blue}$ weakly dominates all other strategies of Blue.

Elimination step 4. $\widehat{p}_2^{red} \not\approx \widehat{p}_1^{blue}$ weakly dominates all other strategies of Red.

The strategy profile surviving iterative elimination is an inefficient equilibrium because Red misreports in round 2 with probability 50%.

D.3 Proof of Lemma 2

Lemma 3 shows that the main part of the expected payment can be expressed by δ, and shows the reason why it incentivizes the agents to deviate in a synchronous way with alternating signs. Lemma 2 is a direct consequence of it (using the fact that $\mathbb{E}(p_T^{red} \cdot p_T^{blue}) = p_0^{red} \cdot p_0^{blue}$).

Lemma 3. *In Example 5.1,*

$$\mathbb{E}\Big(\sum_{t=1}^{T} \gamma_t^{blue}\Big) = 100 \cdot \mathbb{E}\bigg(\Big(\sum_{t=1}^{T-1} \big(\delta_t^{red} - \delta_{t-1}^{red}\big)\delta_t^{blue}\Big) - \delta_{T-1}^{red}\delta_T^{blue} + p_0^{red}p_0^{blue} - p_T^{red}\widehat{p}_T^{blue}\bigg) \quad (75)$$

$$\mathbb{E}\Big(\sum_{t=1}^{T} \gamma_t^{red}\Big) = 100 \cdot \mathbb{E}\bigg(\Big(\sum_{t=1}^{T-1} \big(\delta_t^{blue} - \delta_{t-1}^{blue}\big)\delta_t^{red}\Big) - \delta_{T-1}^{blue}\delta_T^{red} + p_0^{blue}p_0^{red} - p_T^{red}\widehat{p}_T^{red}\bigg) \quad (76)$$

These are true even if players can observe the past types of each other. We assume only that $\widehat{p}_T^{blue}$ is independent of p_T^{red} conditional on the history until round $T-1$, and vice versa.

Proof. The martingale property of probabilities implies that $\mathbb{E}(p_t^{blue}) = p_{t-1}^{blue}$ and $\mathbb{E}(p_t^{red}) = p_{t-1}^{red}$ for every $t > 0$. Furthermore, as these martingales for Blue and Red are independent, $\mathbb{E}(p_{t_1}^{blue} \cdot p_{t_2}^{red}) = p_0^{blue} \cdot p_0^{red}$ for every t_1 and t_2. Similarly, $p_T^{red} - p_{T-1}^{red}$ is independent of $\widehat{p}_{T-1}^{blue}$, and we assumed that it is also independent of $\widehat{p}_T^{blue}$ and vice versa. These imply the following.

$$\mathbb{E}\Big(\big(p_t^{red} - p_{t-1}^{red}\big) \cdot \widehat{p}_{t-1}^{blue}\Big) = \mathbb{E}\Big(\big(p_t^{red} - p_{t-1}^{red}\big) \cdot \widehat{p}_t^{blue}\Big) = \mathbb{E}\Big(\big(p_t^{red} - p_{t-1}^{red}\big) \cdot \delta_{t-1}^{blue}\Big) = 0 \tag{77}$$

$$\mathbb{E}\Big(\big(p_t^{blue} - p_{t-1}^{blue}\big) \cdot \widehat{p}_{t-1}^{red}\Big) = \mathbb{E}\Big(\big(p_t^{blue} - p_{t-1}^{blue}\big) \cdot \widehat{p}_t^{red}\Big) = \mathbb{E}\Big(\big(p_t^{blue} - p_{t-1}^{blue}\big) \cdot \delta_{t-1}^{red}\Big) = 0 \tag{78}$$

By adding up (35) for every t and taking expectation, we obtain the following.

$$\mathbb{E}\Big(\sum_{t=1}^{T} \gamma_t^{blue}\Big) = 100 \cdot \sum_{t=1}^{T} \mathbb{E}\Big(\widehat{p}_{t-1}^{red} \cdot \big(\widehat{p}_{t-1}^{blue} - \widehat{p}_t^{blue}\big)\Big)$$

$$= 100 \cdot \sum_{t=1}^{T} \mathbb{E}\Big(\delta_{t-1}^{red} \cdot \big(\delta_{t-1}^{blue} - \delta_t^{blue}\big) + \delta_{t-1}^{red} \cdot \big(p_{t-1}^{blue} - p_t^{blue}\big) + p_{t-1}^{red} \cdot \big(\widehat{p}_{t-1}^{blue} - \widehat{p}_t^{blue}\big)\Big)$$

$$\overset{(78)}{=} 100 \cdot \sum_{t=1}^{T} \mathbb{E}\Big(\delta_{t-1}^{red} \cdot \big(\delta_{t-1}^{blue} - \delta_t^{blue}\big) + p_{t-1}^{red} \cdot \big(\widehat{p}_{t-1}^{blue} - \widehat{p}_t^{blue}\big)\Big)$$

$$= 100 \cdot \sum_{t=1}^{T} \mathbb{E}\Big(\delta_{t-1}^{red} \cdot \delta_{t-1}^{blue} - \delta_{t-1}^{red} \cdot \delta_t^{blue} + p_{t-1}^{red} \cdot \widehat{p}_{t-1}^{blue} - p_{t-1}^{red} \cdot \widehat{p}_t^{blue}\big)\Big)$$

$$\overset{(77)}{=} 100 \cdot \mathbb{E}\Big(\sum_{t=1}^{T-1} \delta_t^{red} \cdot \delta_t^{blue} - \sum_{t=1}^{T} \delta_{t-1}^{red} \cdot \delta_t^{blue} + \sum_{t=1}^{T} \big(p_{t-1}^{red} \cdot \widehat{p}_{t-1}^{blue} - p_t^{red} \cdot \widehat{p}_t^{blue}\big)\Big)$$

$$100 \cdot \mathbb{E}\bigg(\Big(\sum_{t=1}^{T-1} \big(\delta_t^{red} - \delta_{t-1}^{red}\big)\delta_t^{blue}\Big) - \delta_{T-1}^{red}\delta_T^{blue} + p_0^{red} p_0^{blue} - p_T^{red}\widehat{p}_T^{blue}\bigg)$$

This proves (75), and we can get the proof of (76) in an analogous way. □

E AGV is vulnerable to weak dominance

E.1 Setup and AGV

Each agent $i \in N$ privately observes his type $\theta_i \in \Theta_i$ from a fixed initial probability distribution $\mu_i \in \Delta(\Theta_i)$, independently, $\mu = \prod \mu_i$. Then a public decision $x \in X$ is made, and each agent gets a payoff $u(\theta_i, x)$. The utility of an agent is the payoff plus the transfer, namely, $U_i = u(\theta_i, x) + t_i$. If $\forall i \in N,\ \forall \theta_i \in \Theta_i \colon \mu_i(\theta_i) > 0$, then we say that this is a **simple discrete setup**.

AGV chooses a decision policy

$$\chi(\theta) \in \arg\max_{x \in X} \sum_{i \in N} u(\theta_i, x). \tag{79}$$

Then the mechanism asks each agent to report a type $\hat{\theta}_i$ and makes the public decision $\chi(\hat{\theta})$. The payment rule is based on the following.

$$\gamma_i(\hat{\theta}_i) = \mathbb{E}_{\tilde{\theta} \sim \mu}\Big(\sum_{j \in N \setminus \{i\}} u_j\big(\tilde{\theta}_j, \chi(\hat{\theta}_i, \tilde{\theta}_{-i})\big) - u_j\big(\tilde{\theta}_j, \chi(\tilde{\theta})\big)\Big) \tag{80}$$

$\gamma_i(\hat{\theta}_i)$ is paid by the other agents shared equally, so the total transfers are as follows.

$$t_i(\hat{\theta}) = \gamma_i(\hat{\theta}_i) - \frac{1}{|N| - 1} \sum_{j \in N \setminus \{i\}} \gamma_j(\hat{\theta}_j) \tag{81}$$

Each efficient decision policy χ in (79) determines an **AGV implementation**.

E.2 Results

Theorem 6. *If there is a unique efficient AGV implementation in a simple discrete setup, then the truthful strategy profile is not weakly dominated.*

Note. This theorem partially "saves" the AGV mechanism, but this argument works only in a static and not in a dynamic environment.

Proof of Theorem 1. In short, any deviation either does not change the public decision, and hereby always provides the same utility as the truthful report, or it possibly leads to an inefficient public decision, and hereby strictly worse if the others are truthful.

The detailed proof is the following. We can assume that $|X| > 1$. We say that two reports of an agent are decision-equivalent if changing the report from one to the other never changes the public decision. Formally, $\theta_i \overset{i}{\sim} \theta'_i \Leftrightarrow \forall \theta_{-i} \in \Theta_{-i} \colon d(\theta_i, \theta_{-i}) = d(\theta'_i, \theta_{-i})$. We extend this definition for the case when $\theta'_i \in \Delta(\Theta_i)$ is a probability distribution on types: $\theta_i \overset{i}{\sim} \theta'_i$ if and only if θ'_i is supported on types decision-equivalent to θ_i.

The utility of agent i is

$$U_i = u\big(\theta_i, \chi(\hat{\theta})\big) + t_i(\hat{\theta}) = u\big(\theta_i, \chi(\hat{\theta})\big) + \gamma_i(\hat{\theta}_i) - \frac{1}{|N| - 1} \sum_{j \in N \setminus \{i\}} \gamma_j(\hat{\theta}_j) \tag{82}$$

$$= u\big(\theta_i, \chi(\hat\theta)\big) + \mathbb{E}_{\tilde\theta\sim\mu}\Big(\sum_{j\in N\setminus\{i\}} u_j\big(\tilde\theta_j, \chi(\hat\theta_i, \tilde\theta_{-i})\big)\Big) - f(\hat\theta_{-i}) \tag{83}$$

where

$$f(\hat\theta_{-i}) = \mathbb{E}_{\tilde\theta\sim\mu}\Big(u_j\big(\tilde\theta_j, \chi(\tilde\theta)\big)\Big) + \frac{1}{|N|-1}\sum_{j\in N\setminus\{i\}} \gamma_j(\hat\theta_j) \tag{84}$$

Therefore, if the other agents but i are truthful, then

$$\mathbb{E}_{\theta\sim\mu}(U_i \mid \hat\theta_{-i} = \theta_{-i}) = \mathbb{E}_{\theta\sim\mu}\Big(u\big(\theta_i, \chi(\hat\theta_i, \theta_{-i})\big) - f(\theta_{-i})\Big) + \mathbb{E}_{\tilde\theta\sim\mu}\Big(\sum_{j\in N\setminus\{i\}} u_j\big(\tilde\theta_j, \chi(\hat\theta_i, \tilde\theta_{-i})\big)\Big) \tag{85}$$

$$= \mathbb{E}_{\theta\sim\mu}\Big(\sum_{j\in N} u_j\big(\theta_j, \chi(\hat\theta_i, \theta_{-i})\big)\Big) + C_i \tag{86}$$

We show that for any agent i with any type $\theta_i \in \Theta_i$, no stochastic report $\hat\theta_i \in \Delta(\Theta_i)$ weakly dominates reporting the true type θ_i. If $\theta_i \overset{i}{\sim} \hat\theta_i$, then the deviation does not change the payoff nor the transfer, so it is not weakly dominating. Otherwise, there exists $\theta_{-i} \in \Theta_{-i}$ and $\theta'_i \in \mathrm{supp}(\mu_i)$ such that $\chi(\theta_i, \theta_{-i}) \neq \chi(\theta'_i, \theta_{-i})$. If the other agents are truthful, then it has a positive probability that the deviation changes the public decision. The efficient public decision is unique, therefore, this deviation leads to lower expected total utility. (86) shows that it means lower expected utility for i, so it does not weakly dominate the truthful report. □

Theorem 7. *There exists a setup so that in one of its AGV implementations, the truthful strategy of an agent is weakly dominated, moreover, iterative elimination of weakly dominated strategies ends up in an inefficient strategy profile (not depending on the order of elimination).*

Proof. See the example below. □

Conjecture 1. *It is an open question whether there exists a setup so that in all of its AGV implementations, iterative elimination of weakly dominated strategies eliminates all efficient strategy profiles (and the order of elimination does not matter). We suspect that the answer is positive.*

Clues for the construction. Consider an AGV implementation. The critical point is the first time when we eliminate a truthful strategy. At this step, an agent has the type θ'_i, but reports θ''_i instead. It is still a non-eliminated case that only i deviates, so (86) shows that the implied public decision should remain efficient. Now consider the AGV implementation with the modification that whenever this agent i makes this report θ'_i, the same public decision is made as if θ_i'' was reported: $\chi'(\theta_i', \theta_{-i}) = \chi(\theta_i'', \theta_{-i})$. For this AGV implementation, the same elimination process fails at this step.

A similar idea is the following. Consider the decision policy which chooses a uniform random efficient public decision. We say that a reporting $(\theta_i, \hat{\theta}_i)$ (allowing $\theta_i = \hat{\theta}_i$) is safe if it does not allow an inefficient public decision as long as the others are truthful. Consider the game where every agent must report a safe reporting, and then the uniform random efficient decision is made according to the reports, with no transfers. Consider a Nash equilibrium of this game. This provides us a mapping from the type profiles to a probability distribution on efficient public decisions. For the AGV implementation implied by this decision policy, eliminating a truthful strategy profile by iterative elimination of weakly dominated strategies is a key challenge. □

E.3 Proof of Theorem 2

Agents 1 and 2 may benefit from an investment $d \in \{N, S, B\}$ (none, small, big) to be made by agent 3. $\Theta_1 = \Theta_2 = \{L, H\}$ (low, high) chosen with equal probabilities, $|\Theta_3| = 1$, and the payoff table is the following.

$u_i(\theta_i, d_i)$	Agent 1 L or H	Agent 2 L or H	Agent 3
N	0	0	0
S	10 or 20	0	-10
B	20 or 32	5 or 10	-30

We choose the efficient decision policy that we ask agent 3 for the investment only if it is socially strictly beneficial. This policy means the following table.

$d(\theta_1, \theta_2)$	$\theta_2 = L$	$\theta_2 = H$
$\theta_1 = L$	N	N
$\theta_1 = H$	S	B

The AGV transfers $t_i = \gamma_i - \frac{1}{2} \sum_{j \neq i} \gamma_j$ are calculated as follows. For convenience, these transfers are modified by constants.

$$\gamma_1(L) = 0 \tag{87}$$

$$\gamma_1(H) = \frac{(0-10) + (10-30)}{2} = -15 \tag{88}$$

$$\gamma_2(L) = 0 \tag{89}$$

$$\gamma_2(H) = \frac{(32-30) - (20-10)}{2} = -4 \tag{90}$$

$$\gamma_3 = 0 \tag{91}$$

Notice that the report of i changes only γ_i, and therefore, the change in y_i and in γ_i are the same. So for reporting H, agent 1 should pay 15, and agent 2 should pay 4, shared equally between the other two agents.

Elimination step 1. If $\theta_1 = H$, then $\hat{\theta}_1 = H$, because this report costs 15 but benefits a payoff 20 or 32. This implies that $\Pr(\hat{\theta}_1 = H) \geq \frac{1}{2}$.

Elimination step 2. If $\theta_2 = H$, then $\hat{\theta}_2 = H$, because it costs 4 but benefits 10 with a probability at least $\frac{1}{2}$. This implies that $\Pr(\hat{\theta}_2 = H) \geq \frac{1}{2}$.

Elimination step 3. If $\theta_1 = L$, then $\hat{\theta}_1 = H$, because this report costs 15 but benefits an expected payoff at least $\frac{10+20}{2} = 15$, but potentially more if $\Pr(\hat{\theta}_2 = H) > \frac{1}{2}$. This step implies that $\hat{\theta}_1 = H$.

Elimination step 4. If $\theta_2 = L$, then $\hat{\theta}_2 = H$, because this report costs 4 but benefits a payoff of 5. This step implies that $\hat{\theta}_2 = H$.

F Comparison of BTM, DPM, and GUM

We now illustrate the difference in the payment rules in BTM, DPM, and GUM. Consider two sellers, agents 1 and 2, who produce supplementary goods (e.g., bolts and nuts), and a buyer, agent 3. The agents engage in a two-period relationship, $T = 2$. In each period $t \in \{1, 2\}$, the buyer buys quantity $x_t \in [0, 1]$ from each of the two sellers. Before the first period, each seller $i \in \{1, 2\}$ privately observes his random type $\tilde{\theta}^i \sim U[0, 1]$, whose realization θ^i determines his cost function $\frac{1}{2}\theta^i(x_t)^2$ in each period $t \in \{1, 2\}$. The buyer's value per unit of the pair of goods in period $t = 1$ is equal to 1 and, in period $t = 2$, it is a random type $\tilde{\theta}^3 \sim U(0, 1]$ whose realization she privately observes between the two periods.

The efficient (surplus-maximizing) quantities x_t^i are given by:

$$\chi_1(\theta^1, \theta^2) = \frac{1}{\theta^1 + \theta^2} \quad \text{and} \quad \chi_2(\theta^1, \theta^2, \theta^3) = \frac{\theta^3}{\theta^1 + \theta^2}.$$

Note that the first-period trade reveals the realization of the sum of the parameters of the two sellers $\theta_1 + \theta_2$.

In the calculation of transfers, we will be using the following expressions:

$$u^1(\theta) = -\frac{\theta^1}{2} \cdot \left(\chi_1(\theta)^2 + \chi_2(\theta)^2\right) \tag{92}$$

$$= -\frac{\theta^1}{2} \cdot \left(\left(\frac{1}{\theta^1 + \theta^2}\right)^2 + \left(\frac{\theta^3}{\theta^1 + \theta^2}\right)^2\right) = -\frac{\theta^1\left(1 + (\theta^3)^2\right)}{2(\theta^1 + \theta^2)^2}; \tag{93}$$

$$\mathbb{E}\left[u^1(\tilde{\theta})\right] = -\frac{2}{3}\ln(2); \tag{94}$$

$$\mathbb{E}\left[u^1(\theta^1, \tilde{\theta}^{-1})\right] = -\frac{2}{3} \cdot \mathbb{E}\left[\frac{\theta^1}{(\theta^1 + \tilde{\theta}^2)^2}\right] = -\frac{2}{3(1 + \theta^1)}; \tag{95}$$

$$\mathbb{E}\left[u^1(\theta^2, \tilde{\theta}^{-2})\right] = -\frac{2}{3} \cdot \mathbb{E}\left(\frac{\tilde{\theta}^1}{(\tilde{\theta}^1 + \theta^2)^2}\right) = \frac{2}{3} \cdot \left(\frac{1}{(1 + \theta^2)} - \ln\left(\frac{1 + \theta^2}{\theta^2}\right)\right); \tag{96}$$

$$\mathbb{E}\left[u^1(\theta^{-3}, \tilde{\theta}^3)\right] = -\frac{2 \cdot \theta^1}{3(\theta^1 + \theta^2)^2}; \tag{97}$$

$$u^3(\theta) = \chi_1(\theta) + \theta^3 \cdot \chi_2(\theta) = \frac{1}{\theta^1 + \theta^2} + \theta^3 \cdot \frac{\theta^3}{\theta^1 + \theta^2} = \frac{1 + (\theta^3)^2}{\theta^1 + \theta^2}; \tag{98}$$

$$\mathbb{E}\left[u^3(\tilde{\theta})\right] = \frac{8 \cdot \ln(2)}{3}; \tag{99}$$

$$\mathbb{E}\left[u^3(\theta^1, \tilde{\theta}^{-1})\right] = \frac{4}{3} \cdot \ln\left(\frac{1 + \theta^1}{\theta^1}\right). \tag{100}$$

BTM (Athey and Segal, 2013). The payments given by the BTM mechanism are

$$\gamma^1(\theta) = \mathbb{E}\big[u^2(\theta^1, \tilde{\theta}^{-1})\big] - \mathbb{E}\big[u^2(\tilde{\theta})\big] + \mathbb{E}\big[u^3(\theta^1, \tilde{\theta}^{-1})\big] - \mathbb{E}\big[u^3(\tilde{\theta})\big] \tag{101}$$

$$= \mathbb{E}\big[-\frac{1}{2} \cdot \theta^2 \cdot \chi_1(\theta^1, \tilde{\theta}^{-1})^2\big] - \mathbb{E}\big[u^2(-\frac{1}{2} \cdot \theta^2 \cdot \chi_1(\tilde{\theta})^2)\big] \tag{102}$$

$$+ \mathbb{E}\big[\chi_1(\theta^1, \tilde{\theta}^{-1})^2 + \theta^3 \cdot \chi_2(\theta^1, \tilde{\theta}^{-1})\big] - \mathbb{E}\big[\chi_1(\tilde{\theta}) + \theta^3 \cdot \chi_2(\tilde{\theta})\big] \tag{103}$$

$$= \frac{2}{3} \cdot \left(\frac{1}{(1+\theta^1)} - \ln\left(\frac{1+\theta^1}{\theta^1}\right)\right) + \frac{2}{3}\ln(2) + \frac{4}{3} \cdot \ln\left(\frac{1+\theta^1}{\theta^1}\right) - \frac{8}{3}\ln(2) \tag{104}$$

$$= \frac{2}{3} \cdot \left(\frac{1}{(1+\theta^1)} + \ln\left(\frac{1+\theta^1}{\theta^1}\right)\right) - \ln(4); \tag{105}$$

$$\gamma^2(\theta) = \mathbb{E}\big[u^1(\theta^2, \tilde{\theta}^{-2})\big] - \mathbb{E}\big[u^1(\tilde{\theta})\big] + \mathbb{E}\big[u^3(\theta^2, \tilde{\theta}^{-2})\big] - \mathbb{E}\big[u^3(\tilde{\theta})\big] \tag{106}$$

$$= \frac{2}{3} \cdot \left(\frac{1}{(1+\theta^2)} + \ln\left(\frac{1+\theta^2}{\theta^2}\right)\right) - \ln(4); \tag{107}$$

$$\gamma^3(\theta) = u^1(\theta) - \mathbb{E}\big[u^1(\theta^{-3}, \tilde{\theta}^3)\big] + u^2(\theta) - \mathbb{E}\big[u^2(\theta^{-3}, \tilde{\theta}^3)\big] = \frac{1}{6} - \frac{(\theta^3)^2}{2}. \tag{108}$$

The transfers $y^i = \gamma^i - \frac{1}{2}\sum_{j \in N \setminus \{i\}} \gamma^j$ for each agent i become:

$$y^1(\theta) = \frac{1}{3} \cdot \left(\frac{2}{(1+\theta^1)} - \frac{1}{(1+\theta^2)} + 2 \cdot \ln\left(\frac{1+\theta^1}{\theta^1}\right) - \ln\left(\frac{1+\theta^2}{\theta^2}\right)\right) + \frac{(\theta^3)^2}{4} - \ln(2) - \frac{1}{12}; \tag{109}$$

$$y^2(\theta) = \frac{1}{3} \cdot \left(-\frac{1}{(1+\theta^1)} + \frac{2}{(1+\theta^2)} - \ln\left(\frac{1+\theta^1}{\theta^1}\right) + 2 \cdot \ln\left(\frac{1+\theta^2}{\theta^2}\right)\right) + \frac{(\theta^3)^2}{4} - \ln(2) - \frac{1}{12}; \tag{110}$$

$$y^3(\theta) = -\frac{1}{3} \cdot \left(\frac{1}{(1+\theta^1)} + \frac{1}{(1+\theta^2)} + \ln\left(\frac{1+\theta^1}{\theta^1}\right) + \ln\left(\frac{1+\theta^2}{\theta^2}\right)\right) - \frac{(\theta^3)^2}{2} + \ln(4) + \frac{1}{6}. \tag{111}$$

To guarantee the participation constraints for the sellers, we can adjust the transfers y^i by adding constant terms:

$$y^1(\theta) + c_1, y^2(\theta) + c_2, \text{ and } y^3(\theta) + c_2 \text{ such that } c_1 + c_2 + c_3 = 0. \tag{112}$$

DPM ([Bergemann and Välimäki, 2010]). Transfers are given by:[39]

$$y^1(\theta) = \frac{\theta^1\big(1+(\theta^3)^2\big)}{2(\theta^1+\theta^2)^2} + \frac{2}{3(\theta^1+\theta^2)} + \frac{2}{3}\ln\left(\frac{\theta^2}{2\cdot(1+\theta^2)}\right) + c_1; \tag{113}$$

$$y^2(\theta) = \frac{\theta^2\big(1+(\theta^3)^2\big)}{2(\theta^1+\theta^2)^2} + \frac{2}{3(\theta^1+\theta^2)} + \frac{2}{3}\ln\left(\frac{\theta^1}{2\cdot(1+\theta^1)}\right) + c_2; \tag{114}$$

$$y^3(\theta) = -\frac{1+(\theta^3)^2}{2(\theta^1+\theta^2)} - \frac{2}{3(\theta^1+\theta^2)} + \frac{8}{3}\ln(2) + c_3, \tag{115}$$

where $c_1 + c_2 + c_3 = 0$.

GUM. Suppose the agents are ordered as either $\{1,2,3\}$ or $\{1,3,2\}$, or $\{3,1,2\}$. Then the transfers can be calculated as:

$$y^1(\theta) = \mathbb{E}\big[u^1(\tilde\theta)\big] - u^1(\theta) + \mathbb{E}\big[(u^1+u^2+u^3)(\theta^1,\tilde\theta^2,\tilde\theta^3)\big] - \mathbb{E}\big[(u^1+u^2+u^3)(\tilde\theta^1,\tilde\theta^2,\tilde\theta^3)\big]; \tag{116}$$

$$y^2(\theta) = \mathbb{E}\big[u^2(\tilde\theta)\big] - u^2(\theta) + \mathbb{E}\big[(u^1+u^2+u^3)(\theta^1,\theta^2,\tilde\theta^3)\big] - \mathbb{E}\big[(u^1+u^2+u^3)(\theta^1,\tilde\theta^2,\tilde\theta^3)\big]; \tag{117}$$

$$y^3(\theta) = \mathbb{E}\big[u^3(\tilde\theta)\big] - u^3(\theta) + \mathbb{E}\big[(u^1+u^2+u^3)(\theta^1,\theta^2,\theta^3)\big] - \mathbb{E}\big[(u^1+u^2+u^3)(\theta^1,\theta^2,\tilde\theta^3)\big]. \tag{118}$$

To simplify the expressions, we first calculate:

$$\mathbb{E}\big[(u^1+u^2+u^3)(\tilde\theta^1,\tilde\theta^2,\tilde\theta^3)\big] = \frac{4\ln(2)}{3}; \tag{119}$$

$$\mathbb{E}\big[(u^1+u^2+u^3)(\theta^1,\tilde\theta^2,\tilde\theta^3)\big] = \frac{2}{3}\ln\left(\frac{1+\theta^1}{\theta^1}\right); \tag{120}$$

$$\mathbb{E}\big[(u^1+u^2+u^3)(\theta^1,\theta^2,\tilde\theta^3)\big] = \frac{2}{3(\theta^1+\theta^2)}; \tag{121}$$

$$\mathbb{E}\big[(u^1+u^2+u^3)(\theta^1,\theta^2,\theta^3)\big] = \frac{1+(\theta^3)^2}{2(\theta^1+\theta^2)}. \tag{122}$$

[39]Note that the transfers in the Dynamic Pivot are not budget balanced.

Thus, the transfers simplify to:

$$y^1(\theta) = -\frac{2}{3}\ln(2) + \frac{\theta^1\big(1+(\theta^3)^2\big)}{2(\theta^1+\theta^2)^2} + \frac{2}{3}\ln\left(\frac{1+\theta^1}{\theta^1}\right) - \frac{4}{3}\ln(2) \tag{123}$$

$$= \frac{\theta^1\big(1+(\theta^3)^2\big)}{2(\theta^1+\theta^2)^2} + \frac{2}{3}\ln\left(\frac{1+\theta^1}{8\cdot\theta^1}\right); \tag{124}$$

$$y^2(\theta) = -\frac{2}{3}\ln(2) + \frac{\theta^2\big(1+(\theta^3)^2\big)}{2(\theta^1+\theta^2)^2} + \frac{2}{3(\theta^1+\theta^2)} - \frac{2}{3}\ln\left(\frac{1+\theta^1}{\theta^1}\right) \tag{125}$$

$$= \frac{\theta^2\big(1+(\theta^3)^2\big)}{2(\theta^1+\theta^2)^2} + \frac{2}{3(\theta^1+\theta^2)} + \frac{2}{3}\ln\left(\frac{\theta^1}{2\cdot(1+\theta^1)}\right); \tag{126}$$

$$y^3(\theta) = \frac{8}{3}\ln(2) - \frac{1+(\theta^3)^2}{\theta^1+\theta^2} + \frac{1+(\theta^3)^2}{2(\theta^1+\theta^2)} - \frac{2}{3(\theta^1+\theta^2)} \tag{127}$$

$$= -\frac{1+(\theta^3)^2}{2(\theta^1+\theta^2)} - \frac{2}{3(\theta^1+\theta^2)} + \frac{2}{3}\ln(16). \tag{128}$$

Similarly to the BTM and Dynamic Pivot payments described above, we can adjust the transfers y^i by adding constant terms to guarantee the participation constraints for the sellers:

$$y^1(\theta)+c_1, y^2(\theta)+c_2, \text{ and } y^3(\theta)+c_3 \text{ such that } c_1+c_2+c_3=0. \tag{129}$$

"Symmetric" GUM. As we discuss in Section H, we can make our mechanism "symmetric" by calculating payments for all the possible orderings of the agents and average them. As a result, we will get a symmetric mechanism with payments equal to the Shapley contributions (up to a constant):

$$y^1(\theta) = \frac{\theta^1\big(1+(\theta^3)^2\big)}{2(\theta^1+\theta^2)^2} + \frac{1}{3(\theta^1+\theta^2)} + \frac{1}{3}\ln\left(\frac{(1+\theta^1)\cdot\theta^2}{16\cdot\theta^1(1+\theta^2)}\right) + c_1; \tag{130}$$

$$y^2(\theta) = \frac{\theta^2\big(1+(\theta^3)^2\big)}{2(\theta^1+\theta^2)^2} + \frac{1}{3(\theta^1+\theta^2)} + \frac{1}{3}\ln\left(\frac{(1+\theta^2)\cdot\theta^1}{16\cdot\theta^2(1+\theta^1)}\right) + c_2; \tag{131}$$

$$y^3(\theta) = -\frac{1+(\theta^3)^2}{2(\theta^1+\theta^2)} - \frac{2}{3(\theta^1+\theta^2)} + \frac{2}{3}\ln(16) + c_3, \tag{132}$$

where $c_1+c_2+c_3=0$.

G Collusion in AGV

The Guaranteed Utility Mechanism differs from the Balanced Team Mechanism in two ways. First, the payment rules are different in case of more than $N=2$ agents. Second, the two mechanisms "handle" the same-period reports in a different way. For example, the two-period example in Athey and Segal (2013) has $N = 2$ agents with only one report at each round. As a result, GUM and BTM coincide. However, we can add a "dummy" agent with no action and a constant zero payoff function to highlight the difference between the two mechanisms. In particular, GUM does not change after this addition. However, BTM changes and any two of the three agents can benefit from colluding, as we illustrate using an example below.

Example. Agents 1 and 2 privately observe their types $\theta_1, \theta_2 \in U[-1, 1]$, distributed independently. A public decision $d \in \{-1, +1\}$ is made. The utility of agent $i = 1, 2$ is

$$u_i(\theta_i, d) = \theta_i \cdot d. \tag{133}$$

The AGV mechanism asks for reports $\hat{\theta}_i$ from each agent i, it makes the public decision

$$d = \text{sign}(\hat{\theta}_1 + \hat{\theta}_2),$$

and the transfers are:

$$t_i = \frac{(\hat{\theta}_{-i})^2 - (\hat{\theta}_i)^2}{2}. \tag{134}$$

In this mechanism, truth-telling is a Bayesian Nash equilibrium. However, problems with collusion arise when at least 3 agents are present. We can illustrate the mechanism being prone to collusion by adding a "dummy" agent with neutral preference (constant utility).

Adding a "dummy" agent to the example. Agents 1 and 2 privately receive a type $\theta_1, \theta_2 \in U[-1, 1]$, distributed independently, and agent 3 (a "dummy") has type $\theta_3 = 0$. Therefore, the utility of agent 3 is zero regardless of the public decision $d \in \{-1, +1\}$. The utility of $i \in \{1, 2, 3\}$ remains

$$u_i(\theta_i, d) = \theta_i \cdot d. \tag{135}$$

The AGV mechanism asks for reports $\hat{\theta}_i$ from each agent i and makes a public decision

$$d = \text{sign}(\hat{\theta}_1 + \hat{\theta}_2 + \hat{\theta}_3),$$

and the transfers are given by:

$$t_1 = -\frac{(\hat{\theta}_1)^2}{2} + \frac{(\hat{\theta}_2)^2}{4} + \frac{1}{12}; \tag{136}$$

$$t_2 = \frac{(\hat{\theta}_1)^2}{4} - \frac{(\hat{\theta}_2)^2}{2} + \frac{1}{12}; \tag{137}$$

$$t_3 = \frac{(\hat{\theta}_1)^2 + (\hat{\theta}_2)^2}{4} - \frac{1}{6}. \tag{138}$$

The reason why Agent 3 makes the mechanism prone to collusion is the following. Agent $i \in \{1,2\}$ pays $\frac{\hat{\theta}_i^2}{2} - \frac{1}{6}$ to the others, shared equally between them (including the dummy player). Now we show that any two agents have a benefit from colluding.

Collusion of agents 1 and 2. If agents 1 and 2 agree to report, for example, $\hat{\theta}_i = \varepsilon \cdot \theta_i$ for a small enough ε, then this deviation does not change the decision d, but it increases their total transfer by $(1-\varepsilon^2)\frac{(\hat{\theta}_1)^2+(\hat{\theta}_2)^2}{4}$. This amount is paid by the dummy agent 3.

Collusion of agents 1 and 3. (Analogous to collusion of agents 2 and 3.)

If agent 1 cares about the total transfer to agents 1 and 3, then his cost of reporting a stronger preference decreases from $\frac{(\hat{\theta}_1)^2}{2}$ to $\frac{(\hat{\theta}_1)^2}{2} - \frac{(\hat{\theta}_1)^2}{4} = \frac{(\hat{\theta}_1)^2}{4}$. Consequently, if agent 2 remains truthful, then agent 1 should report $2 \cdot \hat{\theta}$ rounded to $[-1,1]$, namely,

$$\hat{\theta}_1 = \begin{cases} -1 & \text{if} \quad \theta_1 < -1/2 \\ 2 \cdot \theta_1 & \text{if} \quad -1/2 \le \theta_1 \le 1/2 \\ 1 & \text{if} \quad 1/2 < \theta_1 \end{cases} \tag{139}$$

In addition, the truthfulness of Agent 2 is no longer the best response to this collusive behavior, but Agent 2 is better at underreporting. In particular, if the collusion of agents 1 and 3 is common knowledge, then there is an inefficient equilibrium where agent 1 overreports and agent 2 underreports his true preference. Similarly, if agents 1 and 2 both suspect that the other one may have a collusive agreement with dummy agent 3 (with some probability at least), then they both should underreport their types.

BTM: Dynamic AGV. BTM is a dynamic generalization of AGV. Therefore, BTM inherits the aforementioned issue with collusion. Moreover, in a dynamic environment, the issue might be exacerbated. Namely, one agent deviates and incentivizes others to join the deviation. As we show in Section 5, under BTM, iterative elimination of weakly dominated strategies may end up in only strategy profile which is inefficient.

H Connection to Shapley values

GUM and, in particular, our definition of the externality requires the set of the agents $\mathcal{I} = \{1, ..., N\}$ to be ordered.[40] That is, before the mechanism begins, the ordering of the agents in the set $\mathcal{I} = \{1, ..., N\}$ is specified, and all the calculations are performed given that particular ordering. Our mechanism sustains the truthful strategy profile as an equilibrium with any ordering of the agents while retaining the same properties. Therefore, any convex combination of these payment rules also implements the truthful strategy profile. Thus, we can randomly pick the ordering before the mechanism begins.

Another way to make the mechanism "symmetric" is to calculate payments for all the possible orderings of the agents and average them. As a result, we will get a symmetric mechanism with payments equal to the Shapley contributions. That is, when the (simultaneously updated) reports $\hat{\theta}_t^1, \hat{\theta}_t^2, ..., \hat{\theta}_t^N$ change the anticipated payoff $\mathbb{E}[u^i]$, we let $\gamma_t^{i \to j}$ to be the Shapley value of $\hat{\theta}_t^j$ to this change in $\mathbb{E}[u^i]$ (see Shapley (1953)).[41] Since all the agents obtain a guaranteed expected utility for each of the agents' orderings, they get the same guaranteed expected utility for the average of these payment rules.

More formally, every agent's newly updated report creates a coalition function that determines the externality the report induces on each set of agents. In every period t, for each agent i and for each subset of agents $L \subseteq \mathcal{I}$, the updated type vector θ_t^L imposes an externality $\phi_i(L)$ on i, where

$$\phi_i(L) = \mathbb{E}^{\mu(\theta_t^S, \theta_{t-1}^{\mathcal{I}/L})}\left[u^i\right] - \mathbb{E}^{\mu(\theta_{t-1}^{\mathcal{I}})}\left[u^i\right]. \tag{140}$$

In this symmetric mechanism, agent i pays the Shapley values of the function $\phi_i : 2^{\mathcal{I}} \to R$ to the other agents.

Note that this approach requires the reports to be independent conditional on past observations. Thus, we might lose "the extra freedom" that a report $\hat{\theta}_t^i$ can be dependent on the earlier reports in the same round $\hat{\theta}_t^1, \hat{\theta}_t^2, ..., \hat{\theta}_t^{i-1}$ as well as on the corresponding true types $\theta_t^1, \theta_t^2, ..., \theta_t^{i-1}$.

[40] Safronov (2018) discusses the issues related to agents' ordering in the static environment.

[41] There is also a connection between an agent i updated report's marginal contribution in our mechanism and the marginal contribution axiom (see Young (1985) and Pintér (2015)).

References

Athey, S. and K. Bagwell (2008). Collusion with persistent cost shocks. *Econometrica: Journal of the Econometric Society 76*(3), 493–540.

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

Atkeson, A. and J. Lucas, Robert E. (1992). On efficient distribution with private information. *The Review of Economic Studies 59*(3), 427–453.

Ball, I., M. Jackson, and D. Kattwinkel (2022). Comment on Jackson and Sonnenschein (2007) "Overcoming incentive constraints by linking decisions". *Econometrica: Journal of the Econometric Society 90*(4), 3–7.

Battaglini, M. (2005). Long-term contracting with Markovian consumers. *American Economic Review 95*(3), 637–658.

Battaglini, M. and R. Lamba (2019). Optimal dynamic contracting: The first-order approach and beyond. *Theoretical Economics 14*(4), 1435–1482.

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

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

Bergemann, D. and J. Välimäki (2019). Dynamic mechanism design: An introduction. *Journal of Economic Literature 57*(2), 235–74.

Che, Y.-K. and J. Kim (2006). Robustly collusion-proof implementation. *Econometrica: Journal of the Econometric Society 74*(4), 1063–1107.

Che, Y.-K. and J. Kim (2009). Optimal collusion-proof auctions. *Journal of Economic Theory 144*(2), 565–603.

Courty, P. and L. Hao (2000). Sequential screening. *The Review of Economic Studies 67*(4), 697–717.

Cramton, P. C. and T. R. Palfrey (1995). Ratifiable mechanisms: Learning from disagreement. *Games and Economic Behavior 10*(2), 255–283.

Crémer, J. and M. H. Riordan (1985). A sequential solution to the public goods problem. *Econometrica: Journal of the Econometric Society 53*(1), 77–84.

Csóka, E. (2015, June). Efficient teamwork. working paper cs/0602009v98, arXiv.

Csóka, E., A. Rodivilov, and A. Teytelboym (2023). Efficient teamwork. Working paper.

Escobar, J. F. and J. Toikka (2013). Efficiency in games with Markovian private information. *Econometrica: Journal of the Econometric Society 81*(5), 1887–1934.

Fudenberg, D., D. Levine, and E. Maskin (1994). The folk theorem with imperfect public information. *Econometrica: Journal of the Econometric Society 62*(5), 997–1039.

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

Green, J. and J.-J. Laffont (1978). Satisfactory mechanisms for environments with consumption lower bounds. *Journal of Economic Theory 19*(2), 359–375.

Jackson, M. and H. Sonnenschein (2007). Overcoming incentive constraints by linking decisions. *Econometrica: Journal of the Econometric Society 75*(1), 241–257.

Khalil, F., J. Lawarree, and A. Rodivilov (2020). Learning from failures: Optimal contracts for experimentation and production. *Journal of Economic Theory 190*, 105107.

Khalil, F., J. Lawarree, and A. Rodivilov (2024). Rewarding success and failure: moral hazard and adverse selection in strategic experimentation. Working Paper.

Krishna, V. and M. Perry (1998). Efficient mechanism design. Technical report, University Library of Munich, Germany. Available at SSRN 64934.

Laffont, J.-J. and D. Martimort (1997). Collusion under asymmetric information. *Econometrica: Journal of the Econometric Society 65*(4), 875–911.

Laffont, J.-J. and D. Martimort (2000). Mechanism design with collusion and correlation. *Econometrica: Journal of the Econometric Society 68*(2), 309–342.

Lee, J. and H. Sabourian (2011). Efficient repeated implementation. *Econometrica: Journal of the Econometric Society 79*(6), 1967–1994.

Lee, J. and H. Sabourian (2013). Repeated implementation with incomplete information. Working paper.

Liu, H. (2022, December). Collusion-proof dynamic mechanisms. Technical report, University of Michigan.

Mezzetti, C. and L. Renou (2017). Repeated Nash implementation. *Theoretical Economics 12*(1), 249–285.

Miller, D. A. (2011). Robust Collusion with Private Information. *The Review of Economic Studies 79*(2), 778–811.

Pavan, A., I. Segal, and J. Toikka (2014). Dynamic mechanism design: A Myersonian approach. *Econometrica: Journal of the Econometric Society 82*(2), 601–653.

Pintér, M. (2015). Young's axiomatization of the Shapley value: a new proof. *Annals of Operations Research 235*(1), 665–673.

Renou, L. and T. Tomala (2015). Approximate implementation in Markovian environments. *Journal of Economic Theory 159*, 401–442.

Rodivilov, A. (2022). Monitoring innovation. *Games and Economic Behavior 135*, 297–326.

Safronov, M. (2018). Coalition-proof full efficient implementation. *Journal of Economic Theory 177*, 659–677.

Shapley, L. S. (1953). Stochastic games. *Proceedings of the National Academy of Sciences 39*(10), 1095–1100.

Wang, C. (1995). Dynamic insurance with private information and balanced budgets. *The Review of Economic Studies 62*(4), 577–595.

Young, H. P. (1985). Monotonic solutions of cooperative games. *International Journal of Game Theory 14*(2), 65–72.