A Robust Efficient Dynamic Mechanism

Endre Csóka

Alfréd Rényi Institute of Mathematics, Budapest, Hungary

Abstract

Athey and Segal introduced an efficient budget-balanced mechanism for a dynamic stochastic model with quasilinear payoffs and private values, using the solution concept of perfect Bayesian equilibrium (PBE) [3]. We show that this implementation is not robust in multiple senses, especially for at least 3 agents. For example, we will show a generic setup where all efficient strategy profiles can be eliminated by iterative elimination of weakly dominated strategies. Furthermore, this model used strong assumptions about the information of the agents, and the mechanism was not robust to the relaxation of these assumptions. In this paper, we will show a different mechanism that implements efficiency under weaker assumptions and uses the stronger solution concept of "efficient Nash equilibrium with guaranteed expected payoffs".

1 Introduction

The Vickrey–Clarke–Groves (VCG) mechanism [9, 5, 8] established the existence of an efficient, incentive-compatible mechanism for a general class of static mechanism design problems with private values and quasilinear preferences. Subsequently, Arrow [1] and d'Aspremont and Gérard-Varet (AGV) [7] constructed an efficient, incentive-compatible mechanism in which the transfers were also budget-balanced, using the solution concept of Bayesian–Nash equilibrium, under the additional assumption that private information is independent across agents. In dynamic mechanism design problems with private values, Bergemann and Välimäki [4] and Athey and Segal [3] defined dynamic extensions of the VCG and AGV mechanisms.

In this paper, we will refer to the Athey–Segal balanced Team Mechanism. We will show a very generic setup with at least 3 agents where the truthful PBE is degenerate in multiple senses: unstable, and not robust to the idealized information structure. Then we will also show that a different mechanism implements efficiency in a very stable and robust way, including even collusion-resistance. This mechanism coincides with the balanced Team Mechanism only for 2 agents and with no same-round randomness in the private types.

1.1 The setup and the balanced Team Mechanism

We assume that the Reader knows the setup and the balanced Team Mechanism, but we give a short summary about them.

We have a set $N = \{1, 2, ..., n\}$ of agents with fixed initial types $\theta_0^N \in \Theta^N$ that are publicly known even to the designer, so the mechanism depends on the initial types θ_0^N. There are T number of rounds. In each round t, a public decision $x_t \in X$ is made by the designer, and each agent i gets utility $u(x_t, \theta_t^i)$ where $\theta_t^i \in \Theta$ is his current type. Between consecutive rounds, the type of each agent i is changed by a stochastic function of the previous-round

type and the public decision.[1] Formally, each θ_{t+1}^i is chosen from the probability distribution $\mu(x_t, \theta_t^i) \in \Delta(\Theta)$, these randomizations are independent (and μ is fixed as a rule of the game). The agents can report their types, and the designer can prescribe monetary transfers between the agents y^i with $\sum_i y^i = 0$. Every agent i maximizes his expected total utility plus transfer $\mathsf{E}\big(\sum_t u(x_t, \theta_t^i) + y^i\big)$. Our goal is to maximize the total expected utility.

The information of each agent consists of his own type history, the public reports, and the (implied) public decisions. It is strictly assumed that the agents know nothing else. For example, no agent is capable to reveal to other agents any information correlated with his own past.

Definition 1.1 (**balanced Team Mechanism**). *We fix a decision strategy $\xi : \{1, 2, ..., T\} \times \Theta^N \to X$ which would maximize total expected utility. In every round t, we ask each agent to make a report $\widehat{\theta}_t^i$ about his type. We make the public decision $\xi(t, \widehat{\theta}_t^N)$. For each report $\widehat{\theta}_t^i$, the designer calculates how it changes the total of the expected utilities of others, given the (reported) types of all agents in the previous round $\widehat{\theta}_{t-1}^N$. Formally,*

$$\gamma_t^i = \sum_{j \in N \setminus \{i\}} \sum_{t'=t}^T \mathsf{E}\Big(\upsilon_{t'}^{j,\xi}\big(\widehat{\theta_t^i}, \ \widehat{\theta}_{t-1}^{N \setminus \{i\}}\big) - \upsilon_{t'}^{j,\xi}\big(\widehat{\theta}_{t-1}^N\big)\Big), \tag{1}$$

where $\upsilon_{t'}^{j,\xi}$ expresses the utility of j in round t' given the type profile, and assuming truthful reporting strategies and decision strategy ξ. Each agent i gets this signed transfer γ_t^i paid equally by the other agents, namely, the total transfer to i is

$$y^i = \sum_{t=1}^T \Big(\gamma_t^i - \frac{1}{n-1} \sum_{j \in N \setminus \{i\}} \gamma_t^j\Big). \tag{2}$$

2 The balanced Team Mechanism is not robust

In this section, we are pointing out one of the issues with the mechanism, temporarily accepting the idealized information structure of the model. Our simplest example reveals only one aspect of instability, and this setup will have degeneracies. But then we will show more difficult and generic setups which will show even more serious instabilities.

2.1 Our simplest example for instability

Example 2.1. *We have a large number $n \geq 3$ of agents. Two agents have active roles: Blue and Red, and an agent Green has a passive role. The utilities of others are 0, so their only role is paying the transfers γ to Blue and Red. The game consists of $K \geq 2$ rounds. At the end of the last round, Blue and Red both have a final type HIGH or LOW, and there is a public decision to be made: YES or NO. If NO, then the utility of every agent is 0. If YES, then the utility of (Blue, Red, Green) is (1 or 4, 1 or 4, −6). Whether 1 or 4 depends on the last-round type of the player: 1 with LOW or 4 with HIGH. Therefore, the efficient decision is YES if and only if both agents have HIGH type.*

---

[1] We note that the full version of the model and the mechanism includes private decisions. This was included only in the manuscript version of the Athey–Segal paper [2]. As our counterexample works in this restricted model, it works in the full model, as well. Both versions of the paper also allow verifiable public states which we also skipped from our summary because on the one hand, these can be replaced by an extra agent who cannot lie, and on the other hand, we will not use them in our counterexamples.

Note 2.2. *Example 2.1 is essentially a single-round setup, but we consider earlier rounds in which nothing happens except that the types of the agents are updated. We could easily extend the setup to give some sense to the earlier rounds, see Appendix E.*

For $K = 2$ rounds, we show the table of expected payoffs of Blue and Red with three strategies (to be explained shortly) for each of them. The top left cell (with yellow background) is the truthful equilibrium. The strategy in the second row (and column) weakly dominates the first row (and column, respectively).

$\mathsf{E}\big(\text{payoffs}(\text{Blue, Red})\big)$	$\widehat{p}_t^{red} = p_t^{red}$	$\widehat{p}_t^{red} = p_t^{red}$ except: $\widehat{p}_1^{blue} = 0 \Rightarrow \widehat{p}_2^{red} = 1$	$\widehat{p}_1^{red} = 0$; $\widehat{p}_2^{red} = \big($**if** $\widehat{p}_1^{blue} = 0$ **then** 1 **else** $p_2^{red}\big)$
$\widehat{p}_t^{blue} = p_t^{blue}$	(1, 1)	(1, 1)	(1, 1)
$\widehat{p}_t^{blue} = p_t^{blue}$ except: $\widehat{p}_1^{red} = 0 \Rightarrow \widehat{p}_2^{blue} = 1$	(1, 1)	(1, 1)	(2, 2)
$\widehat{p}_1^{blue} = 0$; $\widehat{p}_2^{blue} = \big($**if** $\widehat{p}_1^{red} = 0$ **then** 1 **else** $p_2^{blue}\big)$	(1, 1)	(2, 2)	$\big(3 - \frac{1}{2(n-1)},\ 3 - \frac{1}{2(n-1)}\big)$

Note 2.3. *If we replaced -6 to, say, -3, and therefore, the efficient decision is* YES *if and only if* **either or both** *agents have* HIGH *type, then we would see the same issue.*

This is our simplest but not nearly our strongest example for instability. We will see that the truthful strategy profile is even less stable by having more rounds or relaxing the idealized information structure.

This example is degenerate in multiple senses, but these are not the cause of instability. In Appendix A, we will show a more convincing but more difficult version of this example with a more complete analysis.

2.2   Explanation of the example and generalizations

The balanced Team Mechanism means the following. In round t, Blue and Red have to report about the current probability p_t^{blue} and p_t^{red} (respectively) that his final type will be HIGH. These reports are denoted by $\widehat{p}_t^{blue}$ and $\widehat{p}_t^{red}$, respectively. (In fact, they should report their types, but the mechanism will use only these probabilities.) Then we have the following transfers

$$\gamma_t^{blue} = (4 - 6) \cdot \widehat{p}_{t-1}^{red} \cdot \widehat{p}_t^{blue} - (4 - 6) \cdot \widehat{p}_{t-1}^{red} \cdot \widehat{p}_{t-1}^{blue} = 2 \cdot \widehat{p}_{t-1}^{red} \cdot \big(\widehat{p}_{t-1}^{blue} - \widehat{p}_t^{blue}\big) \tag{3}$$

$$\gamma_t^{red} = (4 - 6) \cdot \widehat{p}_{t-1}^{blue} \cdot \widehat{p}_t^{red} - (4 - 6) \cdot \widehat{p}_{t-1}^{blue} \cdot \widehat{p}_{t-1}^{red} = 2 \cdot \widehat{p}_{t-1}^{blue} \cdot \big(\widehat{p}_{t-1}^{red} - \widehat{p}_t^{red}\big) \tag{4}$$

paid equally by the $n-1$ other agents. This formula is coming from (1) (or see in [3], page 2473, equation (5)) in the following way. If both agents will have HIGH type, then the public decision will be YES, and the total utility of the agents excluding Blue or Red will be $4 + (-6) = -2$. It will happen with probability $p^{blue} \cdot p^{red}$.

From (3) and (4) and $\mathsf{E}(p_{t+1}^i \mid p_t^i) = p_t^i$, the Reader can easily calculate the values in the table. But the intuition behind the issue is the following.

(3) can be interpreted as $\widehat{p}^{blue}$ is a quantity of a good that Blue can buy and sell for the changing unit price of $2 \cdot \widehat{p}_{t-1}^{red}$ in round t. As a consequence, if an agent expects this unit price to increase (or decrease), then he has an incentive to buy (or sell, respectively) as much as he can. Specifically, the truthful strategy of Blue is weakly dominated by its modification that if $\widehat{p}_{k-1}^{red} = 0$, then he always reports $\widehat{p}_k^{blue} = 1$.

First, we show an intuitive reason why the truthful strategy profile is not rational. If in every odd round both agents report $\widehat{p}_{2t+1}^{blue} = \widehat{p}_{2t+1}^{red} = 1$, and in every even round both report $\widehat{p}_{2t}^{blue} = \widehat{p}_{2t}^{red} = 0$, then in every even round both get a transfer of $2 \cdot 1 \cdot (1 - 0) = 2$, and in every odd round both get a transfer $2 \cdot 0 \cdot (0 - 1) = 0$. So this deviation from the truthful strategies is very beneficial for both agents.

This deviation does not require coordination: if either agent starts playing it, then the other agent is incentivized to join. For example, if Blue reports probability $\widehat{p}_t^{blue} = 0$, then it means that the good is free for Red, so Red should buy as much as he can. Therefore, Blue expects his unit price to increase, as well, so he also buys as much as he can. Namely, both agents report probability $\widehat{p}_{t+1}^{blue} = \widehat{p}_{t+1}^{red} = 1$.

We can modify the example so that the agents never report probability 0 before the last round. Namely, if the rules of the game include that, say, $p_1^{blue}, p_1^{red} \geq 0.1$ possibly with equality, then the same argument holds with $\widehat{p}_1^i = 0.1$ instead of $\widehat{p}_1^i = 0$.

From a more general point of view,

the fact that the truthful strategy profile is a PBE strongly relies on their belief that the reports of the other agents are the best

In more general, if the model allows an arbitrarily small probability that Blue will have a bit different expectation of p_t^{red} from the latest report $\widehat{p}_{t-1}^{red}$, then in every PBE, these two agents will report high and low probabilities alternately, and no PBE will remain close to the truthful strategy profile.

We note that if we replaced 6 to 3, then $\widehat{p}_{2t}^{blue} = 0$, $\widehat{p}_{2t}^{red} = 1$, $\widehat{p}_{2t+1}^{blue} = 1$, $\widehat{p}_{2t+1}^{red} = 0$ would be a beneficial deviation in the same sense.

3 Further weaknesses of the balanced Team Mechanism

3.1 Independent types, information

The balanced Team Mechanism uses the assumption of *independent types*. One may fail to realize that this is a much stronger assumption in a dynamic stochastic game than, say, in an auction game.

First of all, notice that "type" can mean two different things. We call them "payoff type" and "full type". Full type is a synonym of information, but payoff type only includes the information of the player that directly affects his payoff. For example, Harsányi described Bayesian games by full types. But in a game with perfect information, the full types of the players do not really make sense because the full types are always the same for every player (except that each player knows his identity in the game). But payoff types may make sense here. In an auction game, the type of a player typically means payoff type, namely, his own valuation of the good(s).

Independent payoff types is a much more reasonable assumption than independent full types. As we show in Appendix C, the Athey–Segal paper uses neither payoff type nor full type. The only understanding we found that makes the results formally correct is the following. Type means full type excluding the public information. Or with an equivalent interpretation, type means payoff type, but they had a hidden assumption that the information of each agent consists only of his payoff type and the public information.

Whichever interpretation we accept, this is still a very strong and practically unrealistic assumption. For example, this assumption implies that an agent cannot reveal anything about the history of his *past types*. We can see in Section 2 that the mechanism has no tolerance for any relaxation of this assumption. Appendix C.4 can help further understand the nature of this issue.

3.2 Dependence on the initial types

The balanced Athey–Segal mechanism is dependent on the initial types of the agents θ_0^N. This assumes that the initial types are publicly known and verifiable by a court. We can interpret this assumption in different ways, but all of them look practically unreasonable, especially in contrast to the assumption of independent types.

For example, if we are interpreting θ_0^i as common knowledge about the probability distribution of the true type θ_1^i, then it means that the designer can exactly define the common prior and no other agent can possibly have any further knowledge about the prior distribution of the types of others. Based on Example 2.1 (in Section 2), we can see that if an agent has an arbitrary small correlated information about the types of others, then there exists no PBE close in any sense to the truthful strategy profile.

We note that the unbalanced Team Mechanism is not dependent of the initial types of the agents θ_0^N, but this difference between the two mechanisms was not pointed out in the paper.

4 How to fix the balanced Team Mechanism

To fix the Athey–Segal mechanism, we need to separate payoff type θ_t^i from information (full type) ϕ_t^i and ϕ_{t+}^i, where ϕ_t^i is the information of the agent when he reports $\widehat{\theta}_t^i$, and ϕ_{t+}^i is the information of the agent just before the next-round state θ_{t+1}^i is chosen by nature.

We define **momentarily private payoff types** as follows. Each payoff type θ_t^i is independent of the joint information (full type) of the other players in the same round $\phi_t^{N \setminus \{i\}}$ (and everything before round t), conditional on x_{t-1} and θ_{t-1}^i. We can fix the mechanism for momentarily private payoff types. (In fact, we only need an even weaker assumption that each agent i can keep his payoff type θ_t^i momentarily secret if he wants to.)

In the balanced Team Mechanism, the designer calculates γ_t^i as a marginal contribution of θ_t^i given the reported types of the agents in the previous round $\widehat{\theta}_{t-1}^N$. But the same proof would work if it was calculated given the same-round reports of others $(\widehat{\theta}_t^{N \setminus \{i\}}, \widehat{\theta}_{t-1}^i)$. Or we could use any combination of the two rules.

We recommend one combination, which is coming from the quasi-dominant equilibrium implementation of a more complex model in [6] (a tendering model with arbitrary private initial types of the agents), specified for our case. Namely, we update the types one by one. Say, we calculate the change γ_t^i given the reported types $(\widehat{\theta}_t^1, \widehat{\theta}_t^2, ..., \widehat{\theta}_t^{i-1}, \widehat{\theta}_t^i, \widehat{\theta}_{t-1}^{i+1}, ..., \widehat{\theta}_{t-1}^n)$. Formally, (1) is replaced by the following.

$$\gamma_t^i = \sum_{j \in N \setminus \{i\}} \sum_{t'=t}^{T} \mathsf{E}\Big(\upsilon_{t'}^{j,\xi}(\widehat{\theta}_t^1, \widehat{\theta}_t^2, ..., \widehat{\theta}_t^{i-1}, \widehat{\theta}_t^i, \widehat{\theta}_{t-1}^{i+1}, ..., \widehat{\theta}_{t-1}^n) - \upsilon_{t'}^{j,\xi}(\widehat{\theta}_t^1, \widehat{\theta}_t^2, ..., \widehat{\theta}_t^{i-1}, \widehat{\theta}_{t-1}^i, \widehat{\theta}_{t-1}^{i+1}, ..., \widehat{\theta}_{t-1}^n) \Big)$$

We also change the rule of how γ_t^i is paid by the other agents. Instead of sharing it equally between the other agents, we say that whoever gets affected positively (or negatively) by a new report pays (or gets) the same transfer. For example, if Blue makes a new report, and it changes the expected total utility of Red by \$15, and of Green by $-\$5$, then *Blue* gets \$15 - \$5 = \$10. With the original transfer rule, each of the $n-1$ other agents pays $\frac{\$10}{n-1}$. With our new rule, Red pays \$15 and Green gets \$5, and the others pay 0. Formally, the payment to i is the following.

$$y_t^i = \sum_{j \in N \setminus \{i\}} \sum_{t'=t}^{T} \mathsf{E}\Big(\upsilon_{t'}^{j,\xi}(\widehat{\theta}_t^1, \widehat{\theta}_t^2, ..., \widehat{\theta}_t^{i-1}, \widehat{\theta}_t^i, \widehat{\theta}_{t-1}^{i+1}, ..., \widehat{\theta}_{t-1}^n) - \upsilon_{t'}^{j,\xi}(\widehat{\theta}_t^1, \widehat{\theta}_t^2, ..., \widehat{\theta}_t^{i-1}, \widehat{\theta}_{t-1}^i, \widehat{\theta}_{t-1}^{i+1}, ..., \widehat{\theta}_{t-1}^n) \Big)$$

$$- \sum_{j \in N \setminus \{i\}} \sum_{t'=t}^{T} \mathsf{E}\Big(\upsilon_{t'}^{i,\xi}(\widehat{\theta}_t^1, \widehat{\theta}_t^2, ..., \widehat{\theta}_t^{j-1}, \widehat{\theta}_t^j, \widehat{\theta}_{t-1}^{j+1}, ..., \widehat{\theta}_{t-1}^n) - \upsilon_{t'}^{i,\xi}(\widehat{\theta}_t^1, \widehat{\theta}_t^2, ..., \widehat{\theta}_t^{j-1}, \widehat{\theta}_{t-1}^j, \widehat{\theta}_{t-1}^{j+1}, ..., \widehat{\theta}_{t-1}^n) \Big)$$

Under this mechanism, the truthful strategy profile is an **efficient Nash equilibrium with guaranteed expected payoffs**. It means the following (with the payoff functions F_i).

- The truthful strategy s_i^* of every agent i guarantees an expected payoff of at least C_i no matter which strategies the other agents choose. Namely, $\forall s_{N \setminus i} \colon \mathsf{E}\big(F_i(s_i^*, s_{N \setminus i})\big) \geq C_i$;

- These guaranteed expected payoffs C_i sum up to the total expected payoffs of all agents with the efficient strategy profile. Namely, $\sum\limits_{i \in N} C_i = \sup\limits_{s_N} \sum\limits_{i \in N} \mathsf{E}\big(F_i(s_N)\big)$.

Notice that it implies collusion-resistance, as well. All these are also true in a model with private decisions. [6]

We note that for 2 players and no same-round chance events, the balanced Team Mechanism coincides with this new mechanism. This is the reason why the example in the Athey–Segal paper (Section 3 in [3]) works perfectly well.

For the (degenerate) case of same-time chance events, the mechanism depends on the ordering of the indexes of the agents. But any convex combination of these mechanisms satisfy the same properties. Therefore, if we want a symmetric mechanism for the agents, then we can average the payment rules for all permutations of the agents. In other words, γ_t^i is the Shapley contribution of the report $\widehat{\theta}_t^i$ among the reports $\widehat{\theta}_t^N$ to the change in the expected total utility of others (by trustful calculation). For Example 2.1, this means

$$\gamma_t^{blue} = 100 \cdot \frac{\widehat{p}_{t-1}^{red} + \widehat{p}_t^{red}}{2} \big(\widehat{p}_{t-1}^{blue} - \widehat{p}_t^{blue}\big),$$

$$\gamma_t^{red} = 100 \cdot \frac{\widehat{p}_{t-1}^{blue} + \widehat{p}_t^{blue}}{2} \big(\widehat{p}_{t-1}^{red} - \widehat{p}_t^{red}\big).$$

If we use this new mechanism, then it resolves the problems shown about Example 2.1. If the two agents play $\widehat{p}_{2t+1}^{blue} = \widehat{p}_{2t+1}^{red} = 1$ and $\widehat{p}_{2t}^{blue} = \widehat{p}_{2t}^{red} = 0$, then they buy and sell the good for always the same unit price of $100 \cdot \frac{0+1}{2} = 100 \cdot \frac{1+0}{2} = 50$.

Note that if we use a continuous-time model instead of the round-by-round model, then for the calculation of γ_t^i, we can use the order of the receiving times of the reports. And hereby same-time reports do not normally happen. In a continuous-time model, momentarily private payoff types only mean that for each update of a payoff type of an agent, this agent can be the first to report about it, and nobody else can observe and report any correlated information any faster.

We have only one weakness that we did not resolve: the initial types are still fixed. This issue cannot be fixed as nicely as the other problems, but this is possible to handle quite well, due to the property of guaranteed expected payoffs. Namely, this mechanism gives an "almost reduction" of the dynamic stochastic problem to a single-round problem. For example in [6], we are discussing a situation where a principal wants to choose some of the competing agents for a dynamic stochastic multi-agent working process, with no prior assumption about the types of the agents.

5   The more general model and the stronger results

We show the model and mechanism in Section 4 not as a comparison to the Athey–Segal paper, but purely alone. Namely, we show the more general model but still with fixed initial types, and how the direct mechanism in [6] applies here. In Appendix H, we will show how the indirect mechanism in [6] can be applied for this setup, and how it extends nicely to the case of nonquasilinear payoffs.

5.1 General notation

Each agent i at each time point m has a **payoff type** $\theta_m^i \in \Theta$ and **information** (or full type) $\phi_m^i \in \mathcal{I}$. **Strategy** is a mapping from information to actions. We always assume that the information $\phi_m^i \in \mathcal{I}$ of each agent i always includes his earlier information $\phi_{m-1}^i \in \mathcal{I}$, his payoff type θ_m^i and the history of "public" actions (or decisions).[2]

In contrast to the standards in theoretical economics, if we do not specify something in the model, then it will mean that it is not specified. In other words, any setup which satisfies the specifications belong to the set of setups we are discussing.

5.2 The extended model

We have a set $N = \{1, 2, ..., n\}$ of agents. There is a finite number of rounds $T = \{1, 2, \ldots k\}$, each round consists of 4 steps (subrounds). Θ denotes the finite set of possible **payoff types**. The initial payoff types $\theta_0^{N_0} \in \Theta^{N_0}$ are fixed, where $N_0 = N \cup \{0\}$ and θ_t^0 is a public type in round t.

Each round $t \in T$ consists of the following steps.

- The planner makes a public decision $x_t^0 \in X$.

- Each agent $i \in N$ makes a decision $x_t^i \in X$.

- For every agent $i \in N_0$, Nature chooses $\theta_t^i \in \Theta$ from a probability distribution $\mu(\theta_{t-1}^i, \theta_{t-1}^0, x_t^0, x_t^i)$ (where $\mu \colon \Theta^2 \times X^2 \to \Delta(\Theta)$ is a given public function) independently from each other and from the state of the game,[3] conditional on these probability distributions.

- Each agent $i \in N$ sends a report $\widehat{\theta_t^i}$.

There might be further actions by the agents or nature (e.g. cheap talk, signaling, type revelation) that may affect the information of the agents (but not their payoff types).

The information of the planner includes the initial payoff types $\theta_0^{N_0}$ and the history of reports and public states $\widehat{\theta_T^{N_0}}$ (up to the current time point). At the end, the planner determines transfers y^i to agent $i \in N$ with $\sum_{i \in N} y^i = 0$.

The utility of each agent $i \in N$ is $u_i = v(\theta_k^i) + y^i$ (for a given $v \colon \Theta \to \mathbb{R}$).

Assumption of momentarily private types. Denote by ϕ_t^i the information of agent i when he reports $\widehat{\theta_t^i}$. For every $i \in N$, θ_t^i must be independent of ϕ_t^{N-i} (and everything before round t) conditional on θ_{t-1}^i. (*Weaker assumption.* If i is using the truthful strategy, then θ_t^i must be independent of $(\phi_t^{N-i}, \phi_{t-1}^i)$ conditional on θ_{t-1}^i.)[4]

We note that in a continuous-time setup, this assumption only means that if an agent reports every chance event (changes in his type) as soon as he can, then the other agents cannot get to know any stochastic information which is not already reported.

---

[2]Formally, for example, ϕ_m^i includes θ_m^i means that there exists a public function $\tau : \mathcal{I} \to \Theta$ such that if agent i has information ϕ_m^i, then his payoff type must be $\theta_m^i = \tau(\phi_m^i)$.

[3]As a further technical assumption, if nature makes any move between this step and the next step in round t, then its distribution must be independent of θ_t^N.

[4]We need to assume that the agent is able to keep his chance event his private information until he reports it. But the point of the weaker version is that we do not need to assume that he is not able to share this information.

5.3 The mechanism

We show a direct mechanism, namely, we always ask the agents to report their types, and we make recommendations for their private decisions. For calculations, the designer will be trustful in the sense that expected utilities will be calculated as if the agents had the same types as they reported, they would always report the truth and they would always make the hidden decision as recommended. We fix an efficient decision policy $\xi : \{1, 2, ..., T\} \times \Theta^{N_0} \to X \times X^N$. In round t, we learn the new types one by one, so $\widehat{\theta}_{t,i}^N = (\widehat{\theta}_t^1, \widehat{\theta}_t^2, ..., \widehat{\theta}_t^{j-1}, \widehat{\theta}_t^j, \widehat{\theta}_{t-1}^{j+1}, ..., \widehat{\theta}_{t-1}^n)$ is our information about the reported types in round t after learning the types of agents 1, 2, ..., j. Let $\Upsilon_{t,j}^i(\theta^N)$ denote the expected total utility of agent i if $\widehat{\theta}_{t,j}^N = \theta^N$ (and with the trustful assumption). Then for each round $t \in \{1, 2, ..., T\}$ and each pair of different agents $(i, j) \in N \times N$, j pays to i a signed transfer of $\Upsilon_{t,i}^j(\theta^N) - \Upsilon_{t,i-1}^j(\theta^N)$.

Note 5.1. *We could ask the agents to report their types at the same time and average the evaluations with every permutation of the agents. But if we use a (more realistic) continuous-time model, then the reports typically arrive one by one anyway.*

In Appendix H, we will sketch the more general non-revelation mechanism in [6] which works importantly differently if the payoffs of the agents are not quasilinear.

5.4 The equilibrium concept and the proof

As u denoted utility in the other sections, in order to avoid ambiguity, let F denote payoff.

Definition 5.2. *In a stochastic dynamic game with a set of players N, suppose that there is a strategy profile $s_*^N \in \mathcal{S}^N$ and constants C^i satisfying the following.*

$$\forall i \in N, \ \forall s^{N-i} \in \mathcal{S}^{N-i}: \qquad \mathsf{E}\big(F^i(s_*^i, s^{N-i})\big) \geq C^i \tag{5}$$

$$\sup_{s^N \in \mathcal{S}^N} \sum_{i \in N} \mathsf{E}\big(F^i(s^N)\big) = \sum_{i \in N} C^i \tag{6}$$

*Then s_*^N is an **efficient Nash equilibrium with guaranteed expected payoffs**.*

Justification. (5) means that i can guarantee himself an expected payoff C^i by playing s_*^i. Each player $i \in N$ has no hope of getting more expected payoff than the maximum possible total expected payoff of all players minus the sum of the guaranteed expected payoffs of the other players. Therefore, i has no hope of getting more expected payoff than

$$\sup_{s^N \in \mathcal{S}^N} \sum_{j \in N} \mathsf{E}\big(F^j(s_N)\big) - \sum_{j \in N \setminus \{i\}} \mathsf{E}\big(F^j(s_*^N)\big) \overset{(6)}{\leq} \sum_{j \in N} C^j - \sum_{j \in N \setminus \{i\}} C^j = C^i.$$

Consequently, each player $i \in N$ has no incentive to deviate from s_*^i, and therefore, we can rightfully say that s_*^N is an equilibrium.

The same argument holds for coalitions: for any $X \subset N$, no joint strategy profile can provide them a higher total expected payoff than

$$\sup_{s^N \in \mathcal{S}^N} \sum_{j \in N} \mathsf{E}\big(F^j(s^N)\big) - \sum_{j \in N \setminus X} \mathsf{E}\big(F^j(s_*^N)\big) \overset{(6)}{\leq} \sum_{j \in N} C^j - \sum_{j \in N \setminus X} C^j = \sum_{i \in X} C^i. \qquad \square$$

Theorem 5.3. *The truthful strategy profile is an efficient Nash equilibrium with guaranteed expected payoffs, with the truthful strategy profile s_*^N and $C^i = \Upsilon_{0,0}^i$.*

Proof. For (5), it is enough to prove that Υ^i plus the total transfers is a martingale if i is truthful because this sum is C^i at the beginning and $u_i(s^i_*, s^{N-i})$ at the end. The martingale property holds when the agent receives his next-round type or the public type is updated because of the martingale property of the expected value. And whenever the other agent j changes Υ^i, he pays to i the signed difference, so the sum is invariant here.

(6) holds because the truthful strategy profile maximizes the expected total payoff. $\qquad\square$

6   Acknowledgement

References

[1] Kenneth J Arrow. The property rights doctrine and demand revelation under incomplete information. In *Economics and Human Welfare*, pages 23–39. Elsevier, 1979.

[2] Susan Athey and Ilya Segal. An efficient dynamic mechanism. manuscript, 2004-2007, `http://www.cs.cmu.edu/~sandholm/cs15-892F13/EfficientDynamic.Athey%20and%20Segal%2007` 2004-2007.

[3] Susan Athey and Ilya Segal. An efficient dynamic mechanism. *Econometrica*, 81(6):2463–2485, 2013.

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

[5] E.H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.

[6] Endre Csóka. Efficient teamwork. *manuscript, Conference on Economic Design, 2017, York; The 9th annual conference of the Israeli chapter of the Game Theory Society, 2017, Haifa; International Conference on Game Theory, Stony Brook, USA, 2015, arXiv preprint cs/0602009*, 2006 (first version).

[7] Claude d'Aspremont and Louis-André Gérard-Varet. Incentives and incomplete information. *Journal of Public Economics*, 11(1):25–45, 1979.

[8] T. Groves. Incentives in teams. *Econometrica: Journal of the Econometric Society*, pages 617–631, 1973.

[9] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.

Appendix

A   A more generic counterexample

We present a longer but more convincing example to show that the issues with the Athey–Segal balanced Team Mechanism were not caused by the degenerations of our previous examples.

In this section, we will specify and slightly modify Example 2.1. Our analysis will be based on Lemma A.1. In order to get convenient values, we replace the utilities of (Blue, Red, Green) to $(84 \text{ or } 104, 84 \text{ or } 104, -204)$, and hereby the constant factor $6 - 4 = 2$ is replaced by $204 - 104 = 100$ in (3) and (4). We use the notation $\delta_t^i = \widehat{p}_t^i - p_t^i$.

Lemma A.1. *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 $t = k$. Then the followings hold.*

$$\mathsf{E}(\gamma^{blue}) = \mathsf{E}\Big(\sum_{t=1}^{k} \gamma_t^{blue}\Big) = 100 \cdot \sum_{t=1}^{k} \mathsf{E}\big(-\delta_{t-1}^{red} \cdot \delta_t^{blue}\big) = 100 \cdot \sum_{t=1}^{\lfloor \frac{k-2}{2} \rfloor} \mathsf{E}\big(-\delta_{2t}^{red} \cdot \delta_{2t+1}^{blue}\big) \quad (7)$$

$$\mathsf{E}(\gamma^{red}) = \mathsf{E}\Big(\sum_{t=1}^{k} \gamma_t^{red}\Big) = 100 \cdot \sum_{t=1}^{k} \mathsf{E}\big(-\delta_{t-1}^{blue} \cdot \delta_t^{red}\big) = 100 \cdot \sum_{t=1}^{\lfloor \frac{k-1}{2} \rfloor} \mathsf{E}\big(-\delta_{2t-1}^{blue} \cdot \delta_{2t}^{red}\big) \quad (8)$$

A.1   The setup

Consider now Example 2.1 with $k = 4$ number of rounds (and $n \geq 3$ players). We define a finite space of types for both players, the types are encoded by

[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 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 utility -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**). (Formally, $X_2 = \{$"b2:20%","b2:80%"$\}$, $X_3 = \{$"r3:10%","r3:90%"$\}$, $X_4 = \{$"b4:0%","b4:100%"$\} \times \{$"r4:0%","r4:100%"$\} \times \{$"YES","NO"$\}$.)

A.2 The analysis

Under the balanced Team Mechanism, these modifications in the setup induce no transfers between the agents. The only effect is that the agents get a huge punishment if they do not report truthfully in the specified rounds. Therefore, by (a dynamic version of) *strict domination*, we eliminate the possibilities of not telling the truth about these 4 extra public decisions (marked with **bold**). 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 A.1 apply. Therefore,

- Blue is maximizing $\mathsf{E}(\gamma^{blue}) - \frac{1}{n-1}\mathsf{E}(\gamma^{red}) = 100 \cdot \mathsf{E}(-\delta_2^{red} \cdot \delta_3^{blue}) - \frac{100}{n-1}\mathsf{E}(-\delta_1^{blue} \cdot \delta_2^{red})$,

- Red is maximizing $\mathsf{E}(\gamma^{red}) - \frac{1}{n-1}\mathsf{E}(\gamma^{blue}) = 100 \cdot \mathsf{E}(-\delta_1^{blue} \cdot \delta_2^{red}) - \frac{100}{n-1}\mathsf{E}(-\delta_2^{red} \cdot \delta_3^{blue})$.

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 to low and high.

Let us start with the last decision $\widehat{p}_3^{blue}$. As (7) and (8) show, it can only affect $\mathsf{E}(\gamma^{blue}) = 100 \cdot \mathsf{E}(-\delta_2^{red} \cdot \delta_3^{blue})$. 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 $\widehat{p}_3^{blue} \nsim \widehat{p}_2^{red}$.

If we want to be more careful with elimination by *weak dominance*, then we can use instead the following argument. $\widehat{p}_3^{blue} \nsim \widehat{p}_2^{red}$ is strictly better for Blue unless if $\widehat{p}_2^{red} = p_2^{red}$. Therefore, in every Nash equilibrium,

$$\mathsf{P}(\widehat{p}_3^{blue} \sim \widehat{p}_2^{red} \nsim p_2^{red}) = 0, \tag{9}$$

or in other words, $\widehat{p}_3^{blue} \nsim \widehat{p}_2^{red}$ whenever $\widehat{p}_3^{blue}$ is relevant.

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 $\widehat{p}_3^{blue} \nsim \widehat{p}_2^{red}$. But we will keep in mind that Blue will have a final decision with the only effect that it can possibly decrease γ^{blue}.

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 $\mathsf{E}(\gamma^{red}) = 100 \cdot \mathsf{E}(-\delta_1^{blue} \cdot \delta_2^{red})$ and $\mathsf{E}(\gamma^{blue}) = 100 \cdot \mathsf{E}(-\delta_2^{red} \cdot \delta_3^{blue})$.

- The effect on $\mathsf{E}(\gamma^{red})$. If $\widehat{p}_1^{blue} \nsim p_1^{blue}$, then $\widehat{p}_2^{red} \nsim \widehat{p}_1^{blue}$ makes $\mathsf{E}(\gamma^{red})$ higher by $100 \cdot (0.7 - 0.3) \cdot (0.8 - 0.2) = 24$. If $\widehat{p}_1^{blue} = p_1^{blue}$, then $\mathsf{E}(\gamma^{red}) = 0$ independently of $\widehat{p}_2^{red}$.

- The effect on $\mathsf{E}(\gamma^{blue})$. If $p_3^{blue} \sim p_2^{red}$, then $\widehat{p}_2^{red}$ has no effect on $\mathsf{E}(\gamma^{blue})$. But if $p_3^{blue} \nsim p_2^{red}$, then $\widehat{p}_2^{red} \nsim p_2^{red}$ increases $\mathsf{E}(\gamma^{blue})$ by $100 \cdot (0.8 - 0.2) \cdot (0.9 - 0.1) = 48$.

It implies that 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 $\widehat{p}_1^{blue} \nsim p_1^{blue}$.

More formally, assume by contradiction that in any Nash equilibrium, $\varepsilon = \mathsf{P}(\widehat{p}_2^{red} \sim \widehat{p}_1^{blue} \nsim p_2^{red}) > 0$. (9) implies that $0 < \varepsilon = \mathsf{P}(\widehat{p}_2^{red} \sim \widehat{p}_1^{blue} \nsim p_2^{red}) = \mathsf{P}(\widehat{p}_2^{red} \sim \widehat{p}_1^{blue} \nsim p_2^{red} \sim \widehat{p}_3^{blue})$. If Red chose $\widehat{p}_2^{red} = p_2^{red}$ instead, then it would weakly increase $\mathsf{E}(\gamma^{red})$, and strictly decrease $\mathsf{E}(\gamma^{blue})$ by $48 \cdot \varepsilon$ independently of $\widehat{p}_3^{blue}$. Therefore, this deviation would strictly increase the payoff of Red, which is a contradiction.

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 $\mathsf{E}(\gamma^{red})$ and $\mathsf{E}(\gamma^{blue})$. 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, $\mathsf{E}(\gamma^{blue})$ is 48 or 0, and the probabilities depend on $\widehat{p}_1^{blue}$, and also $\mathsf{P}(\delta_2^{red} \neq 0) = 1/2$ 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,

$$\mathsf{P}(\gamma^{blue} = 48) = \mathsf{P}(\delta_3^{blue} \neq 0) = \mathsf{P}(p_3^{blue} = 10\% \mid p_1^{blue} = 30\%) = 3/4.$$

With $\widehat{p}_1^{blue} = p_1^{blue} = 30\%$, it would be $P(\gamma^{blue} > 0) = \mathsf{P}(p_3^{blue} = 90\% \mid p_1^{blue} = 30\%) = 1/4$. 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. (It does not include that Blue is able to decrease γ^{blue} by his last move. We use a normalization factor of $1/3$ for $(\mathsf{E}(\gamma^{blue}), \mathsf{E}(\gamma^{red}))$ in order to have smaller integers.)

$\frac{1}{3}(\mathsf{E}(\gamma^{blue}),\ \mathsf{E}(\gamma^{red}))$	$\widehat{p}_2^{red} = p_2^{red}$	$\widehat{p}_2^{red} \nsim \widehat{p}_1^{blue}$
$\widehat{p}_1^{blue} = p_1^{blue}$	(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 $\widehat{p}_2^{red} \nsim \widehat{p}_1^{blue}$. We can conclude it by iterative elimination of dominated strategies, or we could just use our intuitive understanding of rationality.

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 $(\widehat{p}_1^{blue} = 70\%$ and $p_2^{red} = 20\%)$;

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

Now we get the following 4×4 game. Remember that we got it by iterative elimination of *only strictly dominated strategies* (in a dynamic sense) and then assuming that the last move is $\widehat{p}_3^{blue} \nsim \widehat{p}_2^{red}$ justified by that $\widehat{p}_3^{blue} \sim \widehat{p}_2^{red}$ never increases $\mathsf{E}(\gamma^{blue})$ and never changes $\mathsf{E}(\gamma^{red})$.

$\frac{1}{3}(\mathsf{E}(\gamma^{blue}),\ \mathsf{E}(\gamma^{red}))$	$\widehat{p}_2^{red} = p_2^{red}$	$\widehat{p}_2^{red} \nsim \widehat{p}_1^{blue}$	$\widehat{p}_2^{red}$ 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)

Notice that the truthful strategy profile is also a Nash equilibrium and a PBE: $\widehat{p}_1^{blue} = p_1^{blue}$ and $\widehat{p}_2^{red} = p_2^{red}$, and Blue decreases $\mathsf{E}(\gamma^{blue})$ to constant 0 by $\widehat{p}_3^{blue} = p_3^{blue}$.

A.3 Further important variants of the setup

1. In order to get a counterexample not only for truthfulness but also for efficiency, we can add a public decision to the setup in round 2, where Red gets an additional utility 1 if he reported the truth. It changes the matrix game a bit but otherwise it does not change the entire argument.

2. If we modify the setup so that p_1^{blue} is uniform random from the interval $[30\%, 70\%]$, then it makes the truthful strategy profile even less reasonable. Because if Blue reports $\widehat{p}_1^{blue} = 30\%$ or $\widehat{p}_1^{blue} = 70\%$, then it is even harder for Red having no doubt that Blue was just telling the truth.

3. If Blue has an option to reveal his type to Red (which is a relaxation of the model), then the truthful strategy profile will no longer be subgame-perfect. Because it is easy to see that revealing $\widehat{p}_1^{blue} = 1 - p_1^{blue}$ would strictly incentivize Red to choose $\widehat{p}_2^{red} \approx \widehat{p}_1^{blue}$, which is strictly better for Blue. (If Blue reveals his type to Red, then a subgame starts from round 2. See Appendix D about the relation between PBE and subgame-perfection.)

A.4 Proof of Lemma A.1

Lemma A.2 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 A.1 is a direct consequence of it (using the fact that $\mathsf{E}(p_k^{red} \cdot p_k^{blue}) = p_0^{red} \cdot p_0^{blue}$).

Lemma A.2. *In Example 2.1,*

$$\mathsf{E}\Big(\sum_{t=1}^{k}\gamma_t^{blue}\Big) = 100 \cdot \mathsf{E}\Big(\Big(\sum_{t=1}^{k-1}\big(\delta_t^{red}-\delta_{t-1}^{red}\big)\delta_t^{blue}\Big) - \delta_{k-1}^{red}\delta_k^{blue} + p_0^{red}p_0^{blue} - p_k^{red}\widehat{p}_k^{blue}\Big) \qquad (10)$$

$$\mathsf{E}\Big(\sum_{t=1}^{k}\gamma_t^{red}\Big) = 100 \cdot \mathsf{E}\Big(\Big(\sum_{t=1}^{k-1}\big(\delta_t^{blue}-\delta_{t-1}^{blue}\big)\delta_t^{red}\Big) - \delta_{k-1}^{blue}\delta_k^{red} + p_0^{blue}p_0^{red} - p_k^{red}\widehat{p}_k^{red}\Big) \qquad (11)$$

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

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

$$\mathsf{E}\Big(\big(p_t^{red}-p_{t-1}^{red}\big)\cdot\widehat{p}_{t-1}^{blue}\Big) = \mathsf{E}\Big(\big(p_t^{red}-p_{t-1}^{red}\big)\cdot\widehat{p}_t^{blue}\Big) = \mathsf{E}\Big(\big(p_t^{red}-p_{t-1}^{red}\big)\cdot\delta_{t-1}^{blue}\Big) = 0 \qquad (12)$$

$$\mathsf{E}\Big(\big(p_t^{blue}-p_{t-1}^{blue}\big)\cdot\widehat{p}_{t-1}^{red}\Big) = \mathsf{E}\Big(\big(p_t^{blue}-p_{t-1}^{blue}\big)\cdot\widehat{p}_t^{red}\Big) = \mathsf{E}\Big(\big(p_t^{blue}-p_{t-1}^{blue}\big)\cdot\delta_{t-1}^{red}\Big) = 0 \qquad (13)$$

If add up (3) for every t and we take expectation, then we get the following.

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

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

13

$$\overset{(13)}{=} 100 \cdot \sum_{t=1}^{k} \mathsf{E}\Big(\delta_{t-1}^{red} \cdot \big(\delta_{t-1}^{blue} - \delta_t^{blue} \big) + p_{t-1}^{red} \cdot \big(\widehat{p}_{t-1}^{blue} - \widehat{p}_t^{blue} \big) \Big)$$

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

$$\overset{(12)}{=} 100 \cdot \mathsf{E}\Big(\sum_{t=1}^{k-1} \delta_t^{red} \cdot \delta_t^{blue} - \sum_{t=1}^{k} \delta_{t-1}^{red} \cdot \delta_t^{blue} + \sum_{t=1}^{k} \big(p_{t-1}^{red} \cdot \widehat{p}_{t-1}^{blue} - p_t^{red} \cdot \widehat{p}_t^{blue} \big) \Big)$$

$$100 \cdot \mathsf{E}\Big(\Big(\sum_{t=1}^{k-1} \big(\delta_t^{red} - \delta_{t-1}^{red} \big) \delta_t^{blue} \Big) - \delta_{k-1}^{red} \delta_k^{blue} + p_0^{red} p_0^{blue} - p_k^{red} \widehat{p}_k^{blue} \Big)$$

This proves (10), and we can get the proof of (11) in an analogous way. $\square$

B   Detailed analysis of Example 2.1

We try to analyze Example 2.1 with small numbers K of rounds and assume that the agents are always able to report arbitrary probabilities in $[0,1]$. Keep in mind that the message of the paper applies better to larger K. One of the purposes of this section is to show the reason why we did not give a full characterization of the PBEs in Example 2.1.

If $K = 1$, then PBE only means Nash equilibrium, and we play a variant of the coordination game (or battle of sexes). There are 2 stable and 2 or 3 unstable BPEs depending on p_0^{blue} and p_0^{red}.

1. The agents always report $\widehat{p}_1^{blue} = \widehat{p}_1^{red} = 0$.

2. If $p_0^{blue}, p_0^{red} \leq 84\%$, then another equilibrium is always reporting $\widehat{p}_1^{blue} = \widehat{p}_1^{red} = 1$.

3. The agents always report the truth, namely, $\widehat{p}_1^{blue} = p_1^{blue} \in \{0,1\}$ and $\widehat{p}_1^{red} = p_1^{red} \in \{0,1\}$.

4. If $p_1^{blue} = 0$, then $\widehat{p}_1^{blue} = 0$. If $p_1^{blue} = 1$, then $\mathsf{P}(\widehat{p}_1^{blue} = 1) = \frac{100}{104}$. The same applies for Red.

5. Only for $p_0^{blue}, p_0^{red} \leq 84\%$. If $p_1^{blue} = 1$, then $\widehat{p}_1^{blue} = 1$. $\mathsf{P}\big((p_1^{blue} = 0) \text{ AND } (\widehat{p}_1^{blue} = 1) \big) = \frac{16}{84} p_0^{blue}$. The same applies for Red.

We note that if cheap talk is allowed, then the set of PBEs is richer. It can be even richer with the possibility of some noisy signaling or other intermediate levels of communication.

Even though $K = 1$ used a binary type and a binary report per agent, it already had a number of PBEs. Now it is easier to believe that if $K \geq 2$, where the first signal and the first report are both an arbitrary real number from $[0,1]$, then the set of PBEs is very rich, we are not able to characterize it.

From an applied point of view, there are two reasonable PBEs. One is $\widehat{p}_1^{blue} = \widehat{p}_1^{red} = 0$, $\widehat{p}_2^{blue} = \widehat{p}_2^{red} = 1$. The other one is $\widehat{p}_1^{blue} = \widehat{p}_1^{red} = 1$, $\widehat{p}_2^{blue} = \widehat{p}_2^{red} = 0$. The truthful strategy profile is one of the unstable PBEs, which does not respect weak dominance because after $\widehat{p}_{K-1}^{blue} = 0$, Red is weakly better by reporting $\widehat{p}_K^{red} = 1$.

C   The meaning of "type" in the Athey–Segal paper

This section is a more detailed analysis extending Section 3.1 about the possible meanings of "type" in the Athey-Segal paper, and how it could have been defined.

C.1 If "type" means payoff type

The example of Athey and Segal in [3] (Section 3) may suggest that type means payoff type. However, this is NOT the understanding of that paper. For example, independent payoff types would make no restriction about the information of the players, but it contradicts the following quote.

> "independent types (...) means that, conditional on decisions (...), an agent's private information does not have any effect on the distribution of the current (...) types of other agents"

We note that the word "effect" is a bit misleading here. For example, in their sense, if i knows θ_1^j, then his knowledge "has an effect" on the distribution of θ_t^j.

However, the paper and the results could be modified so that "type" would mean payoff type, and they could have proved the same theorem with essentially the same proof. Because by the same reason as we explained in Appendix D, if we allowed the agents to observe the past of each other to any degree, then formally this would not rule out the "unconvincing" PBE.

C.2 If "type" means full type

In this case, θ_{t+1}^i includes all public reports from the previous round $\widehat{\theta}_t^N$. This contradicts that it is chosen by a stochastic function $\mu(x_t^i, \theta_t^i)$, because $\widehat{\theta}_t^N$ cannot be written as a function of x_t^i and θ_t^i.

C.3 If "type" means full type excluding the listed public information

This understanding looks unnatural, but it is consistent with the paper. However, with this understanding, the independent types assumption still means that this is a very restrictive model. Formally, it does not even allow two agents to observe the same weather because it would contradict the assumption of independent types. To be fair, we know from different arguments that this alone cannot spoil a PBE-implementation. But the model and the PBE-implementation definitely do not allow an agent to reveal any information that is correlated with his past payoff type.

For example, assume that Blue can only have two types, HIGH or LOW, and his type is constant throughout the game. Assume that Blue has LOW type but Red believes it is HIGH. Then whatever Blue tries to reveal, he cannot shake the faith of Red in his wrong belief, even in the weakest sense: Blue cannot induce Red to make a dominating move which gives Red an extra utility 1 if Blue has type LOW, but it makes no difference if his type is HIGH. Specifically, this assumption means that even if Blue reports HIGH and LOW alternately throughout the game, he cannot prevent Red from always being sure that the latest report is the truth.

C.4 A pure example for an "unconvincing" equilibrium

We give a rough definition of an "unconvincing" implementation using the following example.

Example C.1. *Consider the following game with k rounds and with $n \geq 2$ agents. In every round, every agent chooses YES or NO simultaneously. For each round and each agent i, if i chooses YES and at least one other agent chose YES in the current or in any of the previous rounds, then i gets utility 1, otherwise, he gets 0 for that round. Every agent maximizes his (expected) total utility of the k rounds.*

We did not specify the information of the players (observability of past actions, signaling about them, cheap talk, etc.), so in this sense, this example describes a class of games. The most important version is when the decisions are private information. But from a more applied point of view, we should keep an eye on the cases where the decisions are not perfectly private.

If we use our intuitive understanding of rational behavior, then we see that always choosing YES is the best strategy, and choosing NO has only disadvantages. $k = 1$ round is already an interesting case with (at least) two Nash equilibria. $(NO, NO, ..., NO)$ is the other Nash equilibrium which we call the "unconvincing" equilibrium. Always choosing NO remains a PBE with a larger number of rounds $k \geq 2$, with appropriate beliefs. We can say that this PBE is "even less convincing" than with $k = 1$. Because even after an agent i chose YES, and i would like to share the fact that he chose YES, this equilibrium requires every other agent j to keeping believe that i chose NO. (More precisely, this "unconvincing" PBE is formally valid even if the agents observe the decisions of each other, see Appendix D.)

If a mechanism implements a goal in such an equilibrium, then we say that this is an "unconvincing" implementation. Especially if such an issue happens not only for degenerate setups, but the mechanism transforms completely generic setups into degenerate games and implements a goal with an "unconvincing" equilibrium.

D   Is PBE a refinement of subgame-perfect equilibrium?

PBE is a refinement of subgame-perfect equilibrium, but it applies only for proper subgames, meaning that it is common knowledge that they are in this subgame. But if one of the agents i knows it not for sure but only with probability 1, then it is no longer a proper subgame, and subgame-perfection does not apply here.

According to the standard understanding of information and belief, if an observation is a part of the rules, then it is observed for sure. But if we are speaking about a robust model to observability (for example, the Athey–Segal model with dependent types), then an observation may be made with probability 1 but never for sure. And hereby, if we apply the one-shot deviation principle, then we can find a PBE without noticing that the beliefs of the agents are inconsistent with their observations.

Specifically, the "unconvincing" PBE in Example C.1 is formally valid even if the agents partially or perfectly observe the decisions of each other (unless if we add any perfect observation to the *rules* of the game). Because if every agent believes that everybody else chose and will choose NO, then if agent i observes any YES by another agent j, then it is a 0-probability event according to the belief of agent i. Therefore, the Bayesian update rule allows agent i keeping believe that every other agent chose and will choose NO (and agent i believes that his eyes just dazzled when he observed a YES from j), and agent i can believe that this is still a common belief (with probability 1).

E   Less reporting does not fix the problem

Example 2.1 uses reports which do not have any effect on the public decision. One could naively ask whether we can fix the mechanism by always asking to report the smallest necessary information for the efficient decision.

The smallest problem is that this question is not well-defined because there is no canonical definition of the smallest necessary information to be reported for the efficient decision. But the following two arguments show more clearly that these kinds of ideas cannot fix the problem.

Reason 1. We can easily construct an example to show that an agent can make a useless report of his type essentially by claiming a fake reason. Formally, consider an agent and a

round where his type θ does not have to be reported. Let us modify the setup in the following way. The agent receives an extra bit $b \in \{0, 1\}$, and we add an extra public decision $\widehat{\theta}$. If $b = 1$ and $\widehat{\theta} = \theta$, then the agent gets an extra utility ϵ. If $b = 0$, then $\widehat{\theta}$ has no effect.

Clearly, if $b = 1$, then the agent should report his type. But if $b = 0$ and he wants to, then he can still claim that $b = 1$ and hereby he can report his type. This kind of technique works unless if the designer can rule out every possibility that the agent may make use of a public decision that does not affect the others.

Reason 2. Even a tiny reason can make it necessary for a type to be reported. But from a practical point of view, our counterexample in Section 2 has some robustness. It may be useful to go back to Example C.1 and see what if we introduce an arbitrarily small cost for choosing YES. It would make always choosing NO a formally legit PBE in many senses. Mainly because it will no longer be weakly dominated by any other strategy. However, from a more applied point of view, this equilibrium would still be "pretty unconvincing", especially if k is large. Now if we compare it with Example 2.1, then we can see that the situation is very similar. Therefore, we cannot hope that such a modification could really fix the problem.

F   Weaknesses of the unbalanced Team Mechanism

The unbalanced Team Mechanism also has a weakness in that two colluding agents can get as much payoff as they want. We show an example of this issue.

Consider the following setup with two agents. In each round, each agent receives a signal (payoff type) 1000 or -1 independently, with probabilities of 1/2. Then the designer makes a public decision YES or NO. If it is NO, then both agents get utility 0. But if the decision is YES, then both agents get the utility equal to the signal.

Consider the case when an agent receives a signal -1. If he reports 1000 instead, then it costs him 0 or 2, but it provides 1001 or 999 more utility to the other agent. The former amounts (costing 0 providing 1001) apply if the agent expects the other agent to report 1000. This is a strong motivation for collusion, moreover, it does not even require collusion, especially in the infinite-horizon game. For example, reporting 1000 as long as the other agent also did so (since the beginning, or only in the previous round, or in the previous k rounds, any of these versions work) is another PBE. This equilibrium is not efficient, but it provides a higher payoff for the two agents.

G   How could a valid argument lead to an "unconvincing" mechanism?

The issue primarily arises from the known weaknesses of Perfect Bayesian Equilibrium and the standards of non-robust mechanism design.

PBE is a (slightly imperfect) necessary condition of the intuitive term "rationality", but it is not at all a sufficient condition. Therefore, for analyzing games, it is a meaningful plan to find every PBE of the game, and after that, we make an analysis in which we understand which of the PBEs are meaningful for practice. This latter analysis may not be completely well-defined but it may use our intuitive understanding of rationality. It does not sound ideal, but this is the best we can do. Accordingly, if we design a mechanism to implement a goal in a practically meaningful way, then it is not enough at all to check that, under the mechanism, one of the PBEs satisfies the goal.

In addition, the PBE-implementation of the Athey–Segal mechanism used the one-stage deviation principle. It created an even higher risk of implementing the goal with an "unconvincing"

PBE because of the following reason. PBE has a blind spot about off-equilibrium paths. The Bayesian update rule gives no restriction on the belief of an agent after an off-equilibrium action of another agent. Every belief is allowed here not because every belief is rational, but the definition of PBE does not undertake to tell what cannot be a rational belief in this situation. (PBE is so much permissive that it does not always respect subgame-perfection, see Appendix D.) The one-stage deviation principle (with the revelation principle) uses this freedom of beliefs in a PBE to avoid the rationality conditions about the possibility that another player already deviated. In other words, the one-stage deviation principle is functioning also as a hacking technique that shifts some of the rationality conditions into this blind spot of PBE. Applying the one-stage deviation principle to Example C.1 can help to understand this issue better.

A further important weakness of the balanced Team Mechanism is that it strictly assumes an unrealistic information structure, as we discussed it in Section 3.

H   The original mechanism and nonquasilinear payoffs

The mechanism in [6] was designed for a tendering setup where a number of agents are competing for participation in a project owned by another player called the principal. Rejected agents will do nothing and gain nothing. Now the mechanism was the following. Each agent makes an offer for a contract about the payment rule between him and the principal, as a function of the contractible events including communication. If an agent is accepted, then this payment rule will apply. Then the principal chooses a strategy (about which of them to accept and how to communicate with the agents) that maximizes her minimum possible expected payoff, where expectation applies only to the stochastic changes on her own private type and the public type.

When we apply that original mechanism to this model with fixed initial types, then it means that we assume that the principal and the agents reported truthfully about their initial types, and we consider the continuation of the game with the winner agents of this tender and with the principal. If we assume that the agents have quasilinear payoffs (i.e. payoff = utility + payment), and we restrict the set of reports to the potentially truthful ones, then we get the mechanism in Section 5.3.

The original mechanism does not have any implicit assumption of quasilinear payoffs, and accordingly, that version works "better" if some agents have nonquasilinear payoffs (or they have limited responsibility, etc.). It is very difficult to formally compare two mechanisms partially because the efficient outcome is no longer well-defined here. Therefore, we will just argue in an intuitive way that this is a nicer and more natural generalization of the direct mechanism. For example, if a public decision has an effect only on agent i about whether or not to take a huge risk, then even though i might be very risk-averse, the mechanism in Section 5.3 (or the Team Mechanisms) makes the decision only considering the expected gain. (In simple cases, the agent can hack the mechanism by misreporting, but we cannot get an universal solution by individual hackings.) But the original mechanism chooses the best according to the preference of i.

In more detail, assume that the agents have a payoff function $F(\theta^{0,i}_{1,2,\dots,T}, x_{1,2,\dots,T}, y^i)$ which is monotonic and unbounded by y^i, and every agent maximizes the expected value of F. If we restrict the set of reports to the potentially truthful (or recommended) ones with this weaker assumption, then the mechanism will have the following structure.

Given the types, we determine a decision strategy, the payment rule y^i (with $\sum_{i \in N} y^i \equiv 0$) and the expected payoffs f^i satisfying the following. For each agent i, if he reports his types and makes his private decisions truthfully (implying $\theta^{0,i}_{1,2,\dots,T} = \widehat{\theta}^{0,i}_{1,2,\dots,T}$), then

$$\widehat{\theta}^{N\setminus\{i\}}_{1,2,\dots,T} \to \mathsf{E}_{\theta^{0,i}_{1,2,\dots,T}}\Big(F\big(\theta^{0,i}_{1,2,\dots,T},\ x_{1,2,\dots,T}(\widehat{\theta}^{N_0}_{1,2,\dots,T}),\ y^i(\widehat{\theta}^{N_0}_{1,2,\dots,T})\big)\Big) \equiv f^i,$$

where the distribution of $\theta_{1,2,\ldots,T}^{0,i}$ is also affected by $\widehat{\theta}_{1,2,\ldots,T}^{N\setminus\{i\}}$ via $x_{1,2,\ldots,T}$, and f^i is a constant.

These properties define the mechanism by a backward recursion in the sense that we can recursively define the vectors f^N we can achieve. The point is that each time an agent i reports a change in his private type, he compensates the other agents for the effect, making them neutral about this change. Therefore, if we replace the payoff function of i with its expected value before the last change in his private type considering the induced payments, then we get to the same problem with one round less. Reduction by one of the latest decisions is simpler, we just take the union of the sets of f^N with the different decisions.

It may seem that there is an additional problem here with the choice between the different Pareto optimal vectors f^N (at the beginning of the game). But the truth is that this is essentially the same problem as we had and ignored in the case of quasilinear payoffs. Namely, with quasilinear payoffs, we can add an arbitrary additional balanced transfer rule between the agents depending on their initial types, and constant 0 is not a fair choice in practice. For example, consider the deterministic setup with two agents and a binary public decision to be made in every round, where Blue always prefers YES by 100 and Red always prefers NO by 101. Clearly, every decision should be NO, but any applied mechanism would probably ask Red to pay some compensation for Blue. So the only special property of the case of quasilinear payoffs is that the problems with efficiency and fairness can be split *additively* into two independent problems because the Pareto optimal expected payoffs always form a hyperplane.

We note that in the setup in [6] with competing agents, this specialty does not matter. Essentially because the outside option is a reference point for payoffs or because each agent i submits an offer for a contract including a required expected payoff f^i.